



JunosE™ Software for E Series™ Broadband Services Routers

Broadband Access Configuration Guide

Release

16.1.x



Modified: 2015-08-05

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

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

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

Products made or sold by Juniper Networks or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.

JunosE™ Software for E Series™ Broadband Services Routers Broadband Access Configuration Guide
Release 16.1.x
Copyright © 2015, Juniper Networks, Inc.
All rights reserved.

Revision History
August 2015—FRS JunosE 16.1.x

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Abbreviated Table of Contents

	About the Documentation	xxxi
Part 1	Managing Remote Access	
Chapter 1	Remote Access Overview	3
Chapter 2	Configuring Remote Access	65
Chapter 3	Monitoring and Troubleshooting Remote Access	111
Part 2	Managing RADIUS and TACACS+	
Chapter 4	Configuring RADIUS Attributes	179
Chapter 5	Configuring RADIUS Dynamic-Request Server	227
Chapter 6	Configuring RADIUS Relay Server	235
Chapter 7	RADIUS Attribute Descriptions	241
Chapter 8	Application Terminate Reasons	265
Chapter 9	Monitoring RADIUS	291
Chapter 10	Configuring TACACS+	305
Chapter 11	Monitoring TACACS+	315
Part 3	Managing L2TP	
Chapter 12	L2TP Overview	321
Chapter 13	Configuring an L2TP LAC	331
Chapter 14	Configuring an L2TP LNS	365
Chapter 15	Configuring L2TP Dial-Out	407
Chapter 16	L2TP Disconnect Cause Codes	421
Chapter 17	Monitoring L2TP and L2TP Dial-Out	427
Part 4	Managing DHCP	
Chapter 18	DHCP Overview	461
Chapter 19	DHCP Local Server Overview	471
Chapter 20	Configuring DHCP Local Server	485
Chapter 21	Configuring DHCP Relay	511
Chapter 22	Configuring the DHCP External Server Application	545
Chapter 23	Monitoring and Troubleshooting DHCP	559

Part 5	Managing the Subscriber Environment	
Chapter 24	Configuring Subscriber Management	605
Chapter 25	Monitoring Subscriber Management	615
Chapter 26	Configuring Subscriber Interfaces	619
Chapter 27	Monitoring Subscriber Interfaces	655
Part 6	Managing Subscriber Services	
Chapter 28	Configuring Service Manager	661
Chapter 29	Monitoring Service Manager	749
Part 7	Index	
	Index	777

Table of Contents

	About the Documentation	xxxiii
	E Series and JunosE Documentation and Release Notes	xxxiii
	Audience	xxxiii
	E Series and JunosE Text and Syntax Conventions	xxxiii
	Obtaining Documentation	xxxv
	Documentation Feedback	xxxv
	Requesting Technical Support	xxxvi
	Self-Help Online Tools and Resources	xxxvi
	Opening a Case with JTAC	xxxvi
Part 1	Managing Remote Access	
Chapter 1	Remote Access Overview	3
	Remote Access Overview	4
	B-RAS Data Flow	5
	Configuring IP Addresses for Remote Clients	5
	AAA Overview	5
	Remote Access Platform Considerations	6
	B-RAS Protocol Support	6
	Remote Access References	6
	Mapping a User Domain to a Virtual Router Overview	7
	Mapping User Requests Without a Valid Domain Name	7
	Mapping User Requests Without a Configured Domain Name	8
	Using DNIS	8
	Redirected Authentication	8
	IP Hinting	9
	Domain Name and Realm Name Overview	9
	Using the Realm Name as the Domain Name	10
	Using Delimiters Other Than @	10
	Using Either the Domain or the Realm as the Domain Name	10
	Specifying the Domain Name or Realm Name Parse Direction	11
	Stripping the Domain Name	11
	Stripping the Domain Name Per Virtual Router	11
	Subscriber Username for RID, COA Requests, and Lawful Intercepts	
	When Strip Domain Is Enabled	12
	Using the Strip Domain Functionality Per Virtual Router When Strip	
	Domain Is Enabled for an AAA Domain Map	12
	Redirected Authentication When Strip Domain Is Enabled	13
	Example: Domain Name and Realm Name	13
	Example: Stripping Domain Name Per Virtual Router for RADIUS Server	
	Authentication	14

Single Name Specification for Users from a Domain Overview	16
RADIUS Authentication and Accounting Servers Configuration Overview	17
Server Access	18
Server Request Processing Limit	18
Authentication and Accounting Methods	19
Supporting Exchange of Extensible Authentication Protocol Messages	20
Immediate Accounting Updates	21
Interim Accounting Updates	21
Duplicate and Broadcast Accounting	21
UDP Checksums	22
SNMP Traps and System Log Messages Overview	22
SNMP Traps	22
System Log Messages	23
Local Authentication Servers Configuration Overview	23
Tunnel Subscriber Authentication Configuration Overview	23
Name Server Addresses Configuration Overview	24
Local Address Servers Configuration Overview	25
Local Address Pool Ranges	26
Local Address Pool Aliases	26
Shared Local Address Pools	26
SNMP Thresholds	27
DHCP Features	28
Domain Name Aliases Overview	28
AAA Profile Configuration Overview	29
RADIUS Route-Download Server for Route Distribution Overview	29
Format of Downloaded Routes	30
Framed-Route (RADIUS attribute 22)	30
Framed-IPv6-Route (RADIUS attribute 99)	30
Cisco AV-Pair (Cisco VSA 26-1)	30
How the Route-Download Server Downloads Routes	31
AAA Logical Line Identifier for Subscriber Tracking Overview	31
How the Router Obtains and Uses the LLID	31
RADIUS Attributes in Preauthentication Request	33
Considerations for Using the LLID	33
VSAs for Dynamic IP Interfaces Overview	34
Traffic Shaping for PPP over ATM Interfaces	36
Mapping Application Terminate Reasons and RADIUS Terminate Codes	
Overview	36
Timeout Configuration Overview	38
Limiting Active Subscribers	39
AAA Failure Notification for RADIUS	39
Configuring AAA Session Timeout	39
Standard RADIUS IPv6 Attributes for IPv6 Neighbor Discovery Router	
Advertisements and DHCPv6 Prefix Delegation Configuration	39
Maximum Number of IPv6 Prefixes Assigned to Clients by Using Only the DHCPv6	
Local Server	40

Maximum Number of IPv6 Prefixes Assigned to Clients by Using Both DHCPv6	
Local Server and Neighbor Discovery Router Advertisements	40
Unique IPv6 Prefix Delegated Per Subscriber	41
Same IPv6 Prefix Delegated to Multiple Subscribers	41
Duplicate IPv6 Prefix Check Overview	42
Duplicate IPv6 Prefix Detection in the AAA User Profile Database Overview	42
Guidelines for Duplicate Address Verification	43
Propagation of LAG Subscriber Information to AAA and RADIUS	45
SRC Client Configuration Overview	47
SRC Client and COPS Terminology	47
Retrieval of DSL Line Rate Information from Access Nodes Overview	50
DHCPv6 Local Address Pools for Allocation of IPv6 Prefixes Overview	51
Example: Delegating the DHCPv6 Prefix	54
Order of Preference in Determining the Local Address Pool for Allocating	
Prefixes	54
Order of Preference in Allocating Prefixes and Assigning DNS Addresses to	
Requesting Routers	55
Sending of Delegated-IPv6-Prefix from the DHCPv6 Local Server to SRC Client	
Overview	56
IPv6 Prefix Allocation Using Neighbor Discovery Router Advertisements from	
IPv6 Address Pools Overview	57
Allocation of Neighbor Discovery Prefixes for IPv6 Subscribers over PPP	
Links	57
Order of Preference in Determining the Local Address Pool for Allocating	
Prefixes for Neighbor Discovery Router Advertisements	57
Order of Preference in Assigning Prefixes when Neighbor Discovery Router	
Advertisements are Configured on an Interface	58
Guidelines for Allocating Neighbor Discovery Prefixes Using IPv6 Address	
Pools	58
Understanding IPCP and IPv6CP Negotiations for IPv4 and IPv6 Clients Based	
on RADIUS-Returned Attributes	60
Sending the Acct-Start and Acct-Update Messages After NCP Negotiation Is	
Completed Overview	63
Chapter 2	
Configuring Remote Access	65
Remote Access Configuration Tasks	66
Configuring a B-RAS License	67
Configuring AAA Duplicate Accounting	68
Configuring AAA Broadcast Accounting	68
Overriding AAA Accounting NAS Information	68
Collecting Accounting Statistics	69
Configuring RADIUS AAA Servers	69
Configuring SNMP Traps	71
Creating the AAA Local Authentication Environment	73
Creating AAA Local User Databases	73
Adding AAA User Entries to Local User Databases	74
Adding AAA User Entries to Default Local User Databases	74
Configuring AAA User Entries in Local User Databases	75
Assigning a Local User Database to a Virtual Router	76

Enabling Local Authentication on the Virtual Router	77
Example: Configuring AAA Local Authentication	77
Configuring DNS Primary and Secondary NMS	81
Configuring WINS Primary and Secondary NMS	81
Configuring a Local Address Server	81
Creating an IP Interface	83
Configuring Single PPP Clients per ATM Subinterface	83
Configuring Multiple PPP Clients per ATM Subinterface	84
Controlling Access to Domain Names	85
Example: Associating all Subscribers of a PPP Interface with a Specific Domain Name	86
Example: Associating Multiple Domain Names with a Specific Domain Name . .	87
Configuring an AAA Per-Profile Attribute List	88
Configuring Interim, Broadcast, and Policy-Based Accounting in Virtual Router Groups and PPP Profiles	89
Configuring the NAS-Port-Type Attribute Manually	91
Configuring a Service Description for the AAA Profile	91
Configuring the Route-Download Server to Download Routes	92
Configuring the Router to Obtain the LLID for a Subscriber	93
Troubleshooting Subscriber Preauthentication	95
Configuring Custom Mappings for PPP Terminate Reasons	95
Configuring Duplicate IPv6 Prefix Check	96
Configuring Detection of Duplicate IPv6 Prefixes in the AAA User Profile Database	97
Configuring the SRC Client	97
Configuring the Forwarding of COPS Requests to the SRC Server Based on DCM Profiles	100
Configuring the DHCPv6 Local Address Pools	100
Example: Limiting the Number of Prefixes Used by DHCPv6 Clients	103
Example: Using DHCPv6 Local Address Pools for Prefix Delegation over non-PPP Links	104
Configuring IPv6 Neighbor Discovery Local Address Pools	105
Overriding AAA to Perform IPCP and IPv6CP Negotiations Based on RADIUS-Returned Attributes	107
Sending the Acct-Start and Acct-Update Messages After NCP Negotiation Is Completed	108
Chapter 3	
Monitoring and Troubleshooting Remote Access	111
Setting Baselines for Remote Access	113
Setting a Baseline for AAA Statistics	113
Setting a Baseline for AAA Route Downloads	113
Setting a Baseline for COPS Statistics	114
Setting a Baseline for Local Address Pool Statistics	114
Setting a Baseline for RADIUS Statistics	114
Setting the Baseline for SRC Statistics	114
How to Monitor PPP Interfaces	115
Monitoring AAA Accounting Configuration	115
Monitoring AAA Configuration to Send the Acct-Start Message After NCP Negotiation Is Completed	117

Monitoring AAA Configuration to Send the Acct-Update Message After NCP Negotiation Is Completed	117
Monitoring AAA Accounting Default	118
Monitoring the AAA Accounting Interval	118
Monitoring AAA-Specific Virtual Router Groups	118
Monitoring the Default AAA Authentication Method List	119
Monitoring AAA Domain Name Stripping for a Domain Per Virtual Router	119
Monitoring Mapping Between User Domains and Virtual Routers	120
Monitoring Tunnel Subscriber Authentication	122
Monitoring Routing Table Address Lookup	123
Monitoring the AAA Model	123
Monitoring IP Addresses of Primary and Secondary DNS and WINS Name Servers	123
Monitoring AAA Profile Configuration	124
Monitoring Statistics about the RADIUS Route-Download Server	125
Monitoring Routes Downloaded by the RADIUS Route-Download Server	127
Monitoring Chassis-Wide Routes Downloaded by the RADIUS Route-Download Server	129
Monitoring AAA Statistics	131
Monitoring the Number of Active Subscribers Per Port	133
Monitoring the Maximum Number of Active Subscribers Per-Virtual Router	133
Monitoring Session Timeouts	134
Monitoring Interim Accounting for Users on the Virtual Router	134
Monitoring Virtual Router Groups Configured for AAA Broadcast Accounting	134
Monitoring Configuration Information for AAA Local Authentication	135
Monitoring AAA Server Attributes	136
Monitoring the COPS Layer Over SRC Connection	138
Monitoring Statistics About the COPS Layer	140
Monitoring Local Address Pool Aliases	142
Monitoring Local Address Pools	143
Monitoring Local Address Pool Statistics	144
Monitoring Shared Local Address Pools	145
Monitoring the Routing Table	145
Monitoring the B-RAS License	146
Monitoring the RADIUS Server Algorithm	146
Monitoring RADIUS Override Settings	146
Monitoring the RADIUS Rollover Configuration	147
Monitoring RADIUS Server Information	147
Monitoring RADIUS Services Statistics	149
Monitoring RADIUS SNMP Traps	153
Monitoring RADIUS Accounting for L2TP Tunnels	153
Monitoring RADIUS UDP Checksums	153
Monitoring RADIUS Server IP Addresses	154
Monitoring the RADIUS Attribute Used for IPv6 Neighbor Discovery Router Advertisements	154
Monitoring the RADIUS Attribute Used for DHCPv6 Prefix Delegation	154
Monitoring Duplicate IPv6 Prefixes	154
Monitoring Duplicate IPv6 Prefixes in the AAA User Profile Database	155
Monitoring the SRC Client Connection Status	155

Monitoring SRC Client Connection Statistics	158
Monitoring the SRC Client Version Number	159
Monitoring the SRC Option	160
Monitoring Subscriber Information	161
Monitoring Application Terminate Reason Mappings	167
Monitoring IPv6 Local Pools for DHCP Prefix Delegation By All Configured Pools	169
Monitoring IPv6 Local Pools for DHCP Prefix Delegation By Pool Name	170
Monitoring IPv6 Local Pool Statistics for DHCP Prefix Delegation	171
Monitoring IPv6 Local Pools for Neighbor Discovery Router Advertisements for all Configured Pools	172
Monitoring IPv6 Local Pools for Neighbor Discovery Router Advertisements by Pool Name	173
Monitoring IPv6 Local Pool Statistics for Neighbor Discovery Router Advertisements Allocation of Prefixes	174
Monitoring the Status of the Override Feature to Initiate IPCP and IPv6CP Negotiations Based on RADIUS-Returned Attributes	175

Part 2

Chapter 4

Managing RADIUS and TACACS+

Configuring RADIUS Attributes	179
RADIUS Overview	180
RADIUS Services	180
RADIUS Attributes	181
RADIUS Platform Considerations	181
RADIUS References	181
Subscriber AAA Access Messages Overview	182
RADIUS IETF Attributes Supported for Subscriber AAA Access Messages	183
Juniper Networks VSAs Supported for Subscriber AAA Access Messages	186
Processing DNS Addresses from Microsoft RADIUS VSAs for PPP Clients During IPCP	191
Subscriber AAA Accounting Messages Overview	192
RADIUS IETF Attributes Supported for Subscriber AAA Accounting Messages	193
Juniper Networks VSAs Supported for Subscriber AAA Accounting Messages	196
RADIUS IETF Attributes Supported for AAA Tunnel Accounting Messages	200
AAA Access Messages During IPCP Negotiations for Dual-Stack Subscribers	201
Access-Request Messages When an IPv4 Address is Renegotiated	202
Access-Accept Messages When an IPv4 Address is Assigned	202
AAA Accounting Messages During IPCP Negotiations for Dual-Stack Subscribers	203
Interim-Acct Messages When an IPv4 Address is Assigned	203
Interim-Acct Messages When an IPv4 Address is Not Assigned	204
Interim-Acct Messages During a Successful IPCP Renegotiation for IPv4 Addresses	204
Interim-Acct Messages During a Failed IPCP Renegotiation for IPv4 Addresses	204
Interim-Acct Messages When an IPv4 Address is Released	205
DSL Forum VSAs in AAA Access and Accounting Messages Overview	206

	DSL Forum VSAs Supported for AAA Access and Accounting Messages	206
	RADIUS Attributes Supported for CLI AAA Messages	208
	CLI Commands Used to Modify RADIUS Attributes	209
	CLI Commands Used to Configure RADIUS IETF Attributes	209
	CLI Commands Used to Configure Juniper Networks VSAs	213
	CLI Commands Used to Include ANCP-Related Juniper Networks VSAs in Access and Accounting Messages	215
	CLI Commands Used to Include DSL Forum VSAs in Access and Accounting Messages	217
	CLI Commands Used to Include or Exclude Attributes in RADIUS Messages . . .	218
	CLI Commands Used to Ignore Attributes when Receiving Access-Accept Messages	223
	RADIUS Per-Profile Attribute List Configuration Overview	223
	Example: Configuring RADIUS-Specific Attributes	224
Chapter 5	Configuring RADIUS Dynamic-Request Server	227
	RADIUS Dynamic-Request Server Overview	227
	RADIUS Dynamic-Request Server Platform Considerations	228
	RADIUS Dynamic-Request Server References	229
	Understanding RADIUS-Initiated Disconnect	229
	Disconnect Messages	229
	Message Exchange	229
	Supported Error-Cause Codes (RADIUS Attribute 101)	230
	Qualifications for Disconnect	230
	Security/Authentication	231
	Configuring RADIUS-Initiated Disconnect	231
	Understanding RADIUS-Initiated Change of Authorization	232
	Change-of-Authorization Messages	232
	Message Exchange	232
	Supported Error-Cause Codes (RADIUS Attribute 101)	232
	Qualifications for Change of Authorization	233
	Security/Authentication	234
	Configuring RADIUS-Initiated Change of Authorization	234
Chapter 6	Configuring RADIUS Relay Server	235
	Understanding the RADIUS Relay Server	235
	How RADIUS Relay Server Works	236
	Authentication and Addressing	236
	Accounting	237
	Terminating the Wireless Subscriber's Connection	237
	RADIUS Relay Server Platform Considerations	238
	RADIUS Relay Server References	238
	RADIUS Relay Server and the SRC Software	238
	Using the SRC Software for Addressing	238
	Using the SRC Software for Accounting	239
	Configuring RADIUS Relay Server Support	239
Chapter 7	RADIUS Attribute Descriptions	241
	RADIUS IETF Attributes	241
	Juniper Networks VSAs	248

	DSL Forum VSAs	260
	Pass Through RADIUS Attributes	262
	RADIUS Attributes References	262
Chapter 8	Application Terminate Reasons	265
	AAA Terminate Reasons	265
	L2TP Terminate Reasons	266
	PPP Terminate Reasons	283
	RADIUS Client Terminate Reasons	290
Chapter 9	Monitoring RADIUS	291
	Monitoring Override Settings of RADIUS IETF Attributes	291
	Monitoring the NAS-Port-Format RADIUS Attribute	292
	Monitoring the Calling-Station-Id RADIUS Attribute	293
	Monitoring the NAS-Identifier RADIUS Attribute	293
	Monitoring the Format of the Remote-Circuit-ID for RADIUS	294
	Monitoring the Delimiter Character in the Remote-Circuit-ID for RADIUS	294
	Monitoring the Acct-Session-Id RADIUS Attribute	294
	Monitoring the DSL-Port-Type RADIUS Attribute	294
	Monitoring the Connect-Info RADIUS Attribute	295
	Monitoring the NAS-Port-ID RADIUS Attribute	295
	Monitoring Included RADIUS Attributes	295
	Monitoring Ignored RADIUS Attributes	298
	Setting the Baseline for RADIUS Dynamic-Request Server Statistics	298
	Monitoring RADIUS Dynamic-Request Server Statistics	299
	Monitoring the Configuration of the RADIUS Dynamic-Request Server	300
	Setting a Baseline for RADIUS Relay Statistics	301
	Monitoring RADIUS Relay Server Statistics	301
	Monitoring the Configuration of the RADIUS Relay Server	303
	Monitoring the Status of RADIUS Relay UDP Checksums	303
	Monitoring the Status of ICR Partition Accounting	304
Chapter 10	Configuring TACACS+	305
	Understanding TACACS+	305
	AAA Overview	306
	Administrative Login Authentication	306
	Privilege Authentication	307
	Login Authorization	307
	Accounting	307
	TACACS+ Platform Considerations	309
	TACACS+ References	309
	Retry Attempts for Successful TCP Connection Overview	310
	Configuring TACACS+	310
	Configuring TACACS+ Support	310
	Configuring Authentication	311
	Configuring Accounting	312

Chapter 11	Monitoring TACACS+	315
	Setting Baseline TACACS+ Statistics	315
	Monitoring TACACS+ Statistics	315
	Monitoring TACACS+ Information	317
Part 3	Managing L2TP	
Chapter 12	L2TP Overview	321
	L2TP Overview	321
	L2TP Terminology	322
	Implementing L2TP	323
	Sequence of Events on the LAC	323
	Sequence of Events on the LNS	324
	Packet Fragmentation	325
	L2TP Platform Considerations	326
	L2TP Module Requirements	326
	ERX7xx Models, ERX14xx Models, and the ERX310 Router	326
	E120 Router and E320 Router	327
	Sessions and Tunnels Supported	328
	L2TP References	329
Chapter 13	Configuring an L2TP LAC	331
	LAC Configuration Prerequisites	332
	Modifying L2TP LAC Default Settings for Managing Destinations, Tunnels, and Sessions	332
	Generating UDP Checksums in Packets to L2TP Peers	333
	Specifying a Destruct Timeout for L2TP Tunnels and Sessions	334
	Preventing Creation of New Destinations, Tunnels, and Sessions	334
	Preventing Creation of New Destinations, Tunnels, and Sessions on the Router	335
	Preventing Creation of New Tunnels and Sessions at a Destination	335
	Preventing Creation of New Sessions for a Tunnel	335
	Specifying a Drain Timeout for a Disconnected Tunnel	335
	Shutting Down Destinations, Tunnels, and Sessions	336
	Closing and Preventing Existing and New Destinations, Tunnels, and Sessions on the Router	336
	Closing and Preventing Existing and New Tunnels and Sessions for a Destination	336
	Closing and Preventing Existing and New Sessions in a Specific Tunnel	336
	Closing a Specific Session	337
	Specifying the Number of Retransmission Attempts	337
	Controlling the Transmission of L2TP Hello Messages to L2TP Peers in Control Channels	337
	Configuring Calling Number AVP Formats	338
	Calling Number AVP 22 Configuration Tasks	342
	Configuring the Fallback Format	343
	Disabling the Calling Number AVP	346
	Mapping a User Domain Name to an L2TP Tunnel Overview	347

	Mapping User Domain Names to L2TP Tunnels from Domain Map Tunnel Mode	348
	Mapping User Domain Names to L2TP Tunnels from Tunnel Group Tunnel Mode	352
	Transmission of the Subscriber Access Interface Speed to LNS Using the RX Connect-Speed AVP	354
	Configuring the RX Speed on the LAC	355
	Managing the L2TP Destination Lockout Process	357
	Modifying the Lockout Procedure	357
	Verifying That a Locked-Out Destination Is Available	358
	Configuring a Lockout Timeout	358
	Unlocking a Destination that is Currently Locked Out	359
	Starting an Immediate Lockout Test	359
	Managing Address Changes Received from Remote Endpoints	360
	Configuring LAC Tunnel Selection Parameters	361
	Configuring the Failover Between Preference Levels Method	361
	Configuring the Failover Within a Preference Level Method	362
	Configuring the Maximum Sessions per Tunnel	363
	Configuring the Weighted Load Balancing Method	363
Chapter 14	Configuring an L2TP LNS	365
	LNS Configuration Prerequisites	366
	Configuring an LNS	366
	Creating an L2TP Destination Profile	369
	Creating an L2TP Host Profile	370
	Configuring the Maximum Number of LNS Sessions	371
	Configuring Groups for LNS Sessions	372
	Configuring a VR in the L2TP Host Profile	373
	Configuring the RADIUS Connect-Info Attribute on the LNS	374
	Overriding LNS Out-of-Resource Result Codes 4 and 5	375
	Overriding the Result Codes	375
	Displaying the Current Override Setting	375
	Enabling the L2TP Call Disconnect Notification on Local Pool Exhaustion	376
	Selecting Service Modules for LNS Sessions Using MLPPP	377
	Assigning Bundled Group Identifiers	378
	Overriding All Endpoint Discriminators	379
	Enabling Tunnel Switching	379
	Creating Persistent Tunnels	379
	Testing Tunnel Configuration	380
	Managing L2TP Destinations, Tunnels, and Sessions	380
	Configuring Disconnect Cause Information	381
	Generating the Disconnect Cause AVP Globally	381
	Generating the Disconnect Cause AVP with a Host Profile	381
	Enabling RADIUS Accounting for Disconnect Cause	382
	Displaying Disconnect Cause Statistics	382
	Configuring the Receive Window Size	382
	Configuring the Default Receive Window Size	383
	Configuring the Receive Window Size on the LAC	383
	Configuring the Receive Window Size on the LNS	384

Configuring Peer Resynchronization	385
Configuring Peer Resynchronization for L2TP Host Profiles and AAA Domain Map Tunnels	387
Configuring the Global L2TP Peer Resynchronization Method	388
Using RADIUS to Configure Peer Resynchronization	389
Configuring L2TP Tunnel Switch Profiles	389
Applying the L2TP Tunnel Switch Profile	389
Configuration Guidelines	390
Configuring L2TP AVPs for Relay	390
Configuration Tasks	390
Enabling Tunnel Switching on the Router	391
Configuring L2TP Tunnel Switch Profiles	391
Applying L2TP Tunnel Switch Profiles by Using AAA Domain Maps	392
Applying L2TP Tunnel Switch Profiles by Using AAA Tunnel Groups	393
Applying Default L2TP Tunnel Switch Profiles	393
Applying L2TP Tunnel Switch Profiles by Using RADIUS	394
Configuring the Transmit Connect Speed Calculation Method	394
Transmit Connect Speed Calculation Methods	395
Static Layer 2	396
Dynamic Layer 2	396
QoS	396
Actual	397
Transmit Connect Speed Calculation Examples	397
Example 1: L2TP Session over ATM 1483 Interface	397
Example 2: L2TP Session over Ethernet VLAN Interface	397
Transmit Connect Speed Reporting Considerations	398
Session Termination for Dynamic Speed Timeout	398
Advisory Speed Precedence for VLANs over Bridged Ethernet	398
Using AAA Domain Maps to Configure the Transmit Connect Speed Calculation Method	399
Using AAA Tunnel Groups to Configure the Transmit Connect Speed Calculation Method	399
Using AAA Default Tunnel Parameters to Configure the Transmit Connect Speed Calculation Method	400
Using RADIUS to Configure the Transmit Connect Speed Calculation Method	401
Controlling the Handling of L2TP Hello Packets in Fast Path	402
PPP Accounting Statistics	403
Stateful Line Module Switchover for LNS Sessions	405
Chapter 15	
Configuring L2TP Dial-Out	407
L2TP Dial-Out Overview	407
L2TP Dial-Out Terms	408
L2TP Dial-Out Platform Considerations	409
L2TP Dial-Out References	409
L2TP Dial-Out Network Model	409
L2TP Dial-Out Process	410

	L2TP Dial-Out Operational States	411
	Chassis	411
	Virtual Router	411
	Targets	411
	Sessions	412
	L2TP Dial-Out Outgoing Call Setup Details	414
	Access-Request Message	414
	Access-Accept Message	414
	Outgoing Call	415
	Mutual Authentication	415
	Route Installation	416
	Creating a Profile Before Configuring L2TP Dial-Out	416
	Configuring L2TP Dial-Out	417
	Creating an L2TP Dial-Out Session	417
	Specifying the Maximum Timeout Period for Establishing an L2TP Dial-Out Session	417
	Specifying the Duration for an L2TP Dial-Out Session to Remain in Dormant State	418
	Specifying the Maximum Triggers to Buffer for an L2TP Dial-Out Session	418
	Deleting an L2TP Dial-Out Session	418
	Resetting an L2TP Dial-Out Session	419
Chapter 16	L2TP Disconnect Cause Codes	421
	L2TP Disconnect Cause Codes	421
Chapter 17	Monitoring L2TP and L2TP Dial-Out	427
	Monitoring the Mapping for User Domains and Virtual Routers with AAA	427
	Monitoring Configured Tunnel Groups with AAA	430
	Monitoring Configuration of Tunnel Parameters with AAA	432
	Monitoring Global Configuration Status on E Series Routers	433
	Monitoring Detailed Configuration Information for Specified Destinations	435
	Monitoring Locked Out Destinations	437
	Monitoring Configured L2TP Destination Profiles or Host Profiles	438
	Monitoring Configured and Operational Status of all Destinations	440
	Monitoring Statistics on the Cause of a Session Disconnection	441
	Monitoring Detailed Configuration Information about Specified Sessions	442
	Monitoring Configured and Operational Summary Status	443
	Monitoring Configured Switch Profiles on Router	444
	Monitoring Detailed Configuration Information about Specified Tunnels	445
	Monitoring Configured and Operational Status of All Tunnels	448
	Monitoring Chassis-wide Configuration for L2TP Dial-out	449
	Monitoring Status of Dial-out Sessions	454
	Monitoring Dial-out Targets within the Current VR Context	455
	Monitoring Operational Status within the Current VR Context	456

Part 4	Managing DHCP	
Chapter 18	DHCP Overview	461
	DHCP Overview Information	461
	Session and Resource Control Software	462
	DHCP Platform Considerations	463
	DHCP References	463
	Configuring the DHCP Access Model	464
	Configuring DHCP Proxy Clients	464
	Logging DHCP Packet Information	465
	Viewing and Deleting DHCP Client Bindings	466
	DHCP Client Bindings and Duplicate MAC Addresses for Subinterfaces	
	Overview	469
Chapter 19	DHCP Local Server Overview	471
	Embedded DHCP Local Server Overview	471
	DHCP Local Server and Client Configuration	472
	Equal-Access Mode Overview	472
	Local Pool Selection and Address Allocation	473
	The Connection Process	473
	Standalone Mode Overview	474
	Local Pool Selection and Address Allocation	475
	Server Management Table	476
	DHCP Local Server Prerequisites	477
	DHCP Local Server Configuration Tasks	477
	DHCP Unique ID for Clients and Servers Overview	478
	Authentication and Accounting of IPv6 Subscribers Using the DHCPv6 Local	
	Server Overview	480
	Accounting for IPv6 Subscribers with DHCPv6 Local Server Standalone	
	Mode	481
	Interoperation of Authentication of IPv6 Clients and Display of Active Subscriber	
	Information	482
	Limiting the Maximum Number of IPv6 Prefixes Delegated Per Interface by the	
	DHCPv6 Local Server Overview	483
Chapter 20	Configuring DHCP Local Server	485
	Configuring the DHCP Local Server	485
	Basic Configuration of DHCP Local Server	485
	Limiting the Number of IP Addresses Supplied by DHCP Local Server	487
	Excluding IP Addresses from Address Pools	487
	Configuring DHCP Local Server to Support Creation of Dynamic Subscriber	
	Interfaces	488
	Differentiating Between Clients with the Same Client ID or Hardware	
	Address	488
	Logging Out DHCP Local Server Subscribers	489
	Clearing an IP DHCP Local Server Binding	490
	Using SNMP Traps to Monitor DHCP Local Server Events	490

Chapter 21

Using DHCP Local Server Event Logs	491
Configuring DHCP Local Address Pools	491
Basic Configuration of DHCP Local Address Pools	492
Linking Local Address Pools	493
Setting Grace Periods for Address Leases	494
Sending RADIUS Attributes to DHCP Subscribers Overview	495
Dual-Stack Lite Tunnel Name Configuration with RADIUS and DHCPv6	495
PCP Server Name Configuration with RADIUS and DHCP or DHCPv6	496
Configuring AAA Authentication for DHCP Local Server Standalone Mode	497
Configuring AAA Authentication for DHCPv6 Local Server Standalone Mode	500
Configuring the DHCPv6 Local Server	501
Configuring the Type of DHCP Unique ID for DHCPv6 Local Servers	504
Deleting DHCPv6 Client Bindings	505
Configuring the Router to Work with the SRC Software	507
Configuring DHCP Relay	511
DHCP Relay and BOOTP Relay Overview	512
Configuring DHCP Relay Settings	512
Enabling DHCP Relay	513
Removing Access Routes from Routing Tables and NVS	513
Treating All Packets as Originating at Trusted Sources	514
Assigning the Giaddr to Source IP Address	514
Protecting Against Spoofed Giaddr and Relay Agent Option Values	514
Using the Giaddr to Identify the Primary Interface for Dynamic Subscriber Interfaces	515
Using the Broadcast Flag Setting to Control Transmission of DHCP Reply Packets	516
Interaction of DHCP Relay Broadcast Flag with the Layer 2 Unicast Transmission Method	516
Preventing DHCP Relay from Installing Host Routes by Default	518
Example: Preventing Installation of Host Routes	519
Including Relay Agent Option Values in the PPPoE Remote Circuit ID	520
Configuring the Layer 2 Unicast Transmission Method for Reply Packets to DHCP Clients	521
Using Option 60 Strings to Forward Client Traffic to Specific DHCP Servers	522
Configuring the DHCP Relay Option 60 Attribute for Traffic Forwarding	524
Relaying DHCP Packets That Originate from a Cable Modem	525
DHCP Relay Agent Information Option (Option 82) Suboption Values Overview	525
Format of the JunosE Data Field in the Vendor-Specific Suboption for Option 82	527
Configuring the DHCP Relay Agent Option 82 Settings	529
Configuring Relay Agent Option 82 Information	529
Preventing Option 82 Information from Being Stripped from Trusted Client Packets	530
Configuring Relay Agent Information Option (Option 82) Suboption Values	530
Using the set dhcp relay agent sub-option Command to Enable Option 82 Suboption Support	531

	Using the set dhcp relay agent Command to Enable Option 82 Suboption Support	533
	Example: Using DHCP Relay Option 82 to Pass IEEE 802.1p Values to DHCP Servers	535
	Rate of DHCP Client Packets Processed by DHCP Relay Overview	539
	Manually Configuring the Maximum Rate of Client Packets Processed Per Second by DHCP Relay	540
	Configuring the Rate of Client Packets Processed by DHCP Relay	540
	Configuring DHCP Relay Proxy	540
	Enabling DHCP Relay Proxy	541
	Use the First Offer from a DHCP Server	541
	Set a Timeout for DHCP Client Renewal Messages	541
	Managing Host Routes Using DHCP Relay Proxy	542
	Selecting the DHCP Server Response	542
	Behavior for Bound Clients and Address Renewals	543
Chapter 22	Configuring the DHCP External Server Application	545
	DHCP External Server Overview	545
	Preservation of Dynamic Subscriber Interfaces with DHCP External Server Overview	547
	DHCP External Server Identification of Clients with Duplicate MAC Addresses Overview	548
	Configuration Guidelines for Using Duplicate MAC Mode	549
	Restrictions for Using Duplicate MAC Mode to Manage Clients	549
	DHCP External Server Configuration Requirements	550
	Enabling and Disabling the DHCP External Server Application	550
	Monitoring DHCP Traffic Between Remote Clients and DHCP Servers	551
	Synchronizing the DHCP External Application and the Router	551
	Configuring Interoperation with Ethernet DSLAMs	551
	Configuring the DHCP External Server to Support the Creation of Dynamic Subscriber Interfaces	552
	Configuring DHCP External Server to Control Preservation of Dynamic Subscriber Interfaces	553
	Configuring Dynamic Subscriber Interfaces for Interoperation with DHCP Relay and DHCP Relay Proxy	554
	Deleting Clients from a Virtual Router's DHCP Binding Table	556
	Configuring DHCP External Server to Uniquely Identify Clients with Duplicate MAC Addresses	557
	Configuring DHCP External Server to Re-Authenticate Auto-Detected Dynamic Subscriber Interfaces	558
Chapter 23	Monitoring and Troubleshooting DHCP	559
	Setting Baselines for DHCP Statistics	560
	Setting a Baseline for DHCP Relay and Relay Proxy	560
	Setting a Baseline for DHCP Proxy Server Statistics	560
	Setting a Baseline for DHCP External Server Statistics	561
	Setting a Baseline for DHCP Local Server Statistics	561
	Monitoring Addresses Excluded from DHCP Local Server Use	561
	Monitoring DHCP Bindings	562
	Monitoring DHCP Binding Information	562

Monitoring DHCP Binding Count Information	565
Monitoring DHCP Binding Host Information	567
Monitoring DHCP Bindings (Displaying IP Address-to-MAC Address Bindings)	570
Monitoring DHCP Bindings (Displaying DHCP Bindings Based on Binding ID) . .	570
Monitoring DHCP Bindings (Local Server Binding Information)	571
Monitoring DHCP External Server Configuration Information	572
Monitoring DHCP External Server Statistics	573
Monitoring DHCP External Server Duplicate MAC Address Setting	574
Monitoring DHCP Local Address Pools	575
Monitoring DHCP Local Server Authentication Information	577
Monitoring DHCP Local Server Configuration	578
Monitoring DHCP Local Server Leases	579
Monitoring DHCP Local Server Statistics	580
Monitoring DHCP Option 60 Information	583
Monitoring DHCP Packet Capture Settings	584
Monitoring DHCP Relay Configuration Information	585
Monitoring DHCP Relay Proxy Statistics	586
Monitoring DHCP Relay Statistics	588
Monitoring DHCP Server and DHCP Relay Agent Statistics	591
Monitoring DHCP Server and Proxy Client Information	592
Monitoring DHCPv6 Local Server Binding Information	593
Monitoring DHCPv6 Local Server DNS Search Lists	594
Monitoring DHCPv6 Local Server DNS Servers	594
Monitoring DHCPv6 Local Server Prefix Lifetime	595
Monitoring DHCPv6 Local Server Statistics	595
Monitoring DHCPv6 Local Server Authentication Information	596
Monitoring the Maximum Number of Available Leases	597
Monitoring Static IP Address and MAC Address Pairs Supplied by DHCP Local Server	599
Monitoring Status of DHCP Applications	599
Monitoring DHCP Proxy Client Bindings	600
Monitoring the Maximum Number of IPv6 Prefixes Delegated Per Interface by the DHCPv6 Local Server	600
Monitoring the DHCPv6 Local Server to Send Delegated-IPv6-Prefix to SRC Client	602

Part 5

Chapter 24

Managing the Subscriber Environment

Configuring Subscriber Management	605
Understanding Subscriber Management	605
Subscriber Management Platform Considerations	606
Subscriber Management Attributes	606
Dynamic IP Subscriber Interfaces	607
Subscriber Management Procedure Overview	607
Configuring Subscriber Management with an External DHCP Server	609
Subscriber Management Configuration Examples	610

Chapter 25	Monitoring Subscriber Management	615
	Monitoring IP Service Profiles	615
	Monitoring Active IP Subscribers Created by Subscriber Management	616
Chapter 26	Configuring Subscriber Interfaces	619
	Subscriber Interfaces Overview	619
	Dynamic Interfaces and Dynamic Subscriber Interfaces Overview	620
	Subscriber Interfaces Platform Considerations	622
	Interface Specifiers	622
	Subscriber Interfaces References	623
	Characteristics of Subscriber Interfaces	623
	Relationship to Shared IP Interfaces	624
	Relationship to Primary IP Interfaces	624
	Ethernet Interfaces and VLANs	625
	Moving Interfaces	625
	Preventing IP Spoofing	625
	Routing Protocols	625
	Policies and QoS	625
	Applications for Subscriber Interfaces	626
	Directing Traffic Toward Special Local Content	626
	Differentiating Traffic for VPNs	627
	Dynamic Creation of Subscriber Interfaces Overview	628
	Dynamic Subscriber Interfaces Using DHCP Events Overview	629
	DHCP Local Server and Address Allocation	629
	DHCP External Server and Address Allocation	630
	DHCP Relay Configuration	630
	Supported Configurations	630
	Dynamic Subscriber Interfaces Using Packet Detection Overview	631
	Designating Traffic for the Primary IP Interface	631
	Using Framed Routes	632
	Inheritance of MAC Address Validation State for Dynamic Subscriber Interfaces Overview	632
	How MAC Address Validation State Inheritance Works	633
	Configuration of MAC Address Validation State Inheritance	633
	Verification of MAC Address Validation State Inheritance	634
	Example: Configuring Static Subscriber Interfaces	635
	Example: Using a Destination Address to Demultiplex Traffic	635
	Example: Using a Source Address to Demultiplex Traffic	638
	Example: Configuring Dynamic Subscriber Interfaces	642
	Example: Configuring Dynamic Subscriber Interfaces over Ethernet	642
	Example: Configuring Dynamic Subscriber Interfaces over VLANs	644
	Example: Configuring Dynamic Subscriber Interfaces over Bridged Ethernet	646
	Example: Configuring Dynamic Subscriber Interfaces over GRE Tunnels	648
	Example: Configuring Dynamic Subscriber Interfaces Using Loopback Interfaces	650

Chapter 27	Monitoring Subscriber Interfaces	655
	Monitoring Subscriber Interfaces Overview	655
	Monitoring Subscriber Interfaces	655
	Monitoring Active IP Subscribers Created by Subscriber Management	656
 Part 6	 Managing Subscriber Services	
Chapter 28	Configuring Service Manager	661
	Service Manager Overview	663
	Service Manager Terms and Acronyms	663
	Service Manager Platform Considerations	664
	Service Manager References	665
	Configuring Service Manager	665
	Service Definitions Overview	667
	Creating Service Definitions	668
	Managing Your Service Definitions	671
	Copying a Service Definition Macro File	671
	Installing a Service Definition File	671
	Uninstalling a Service Definition File	672
	Updating an Existing Service Definition File	672
	Referencing Policies in Service Definitions Overview	673
	Referencing QoS Configurations in Service Definitions	673
	Specifying QoS Profiles in Service Definitions	674
	Configuring a QoS Profile for Service Manager	674
	Specifying QoS Profiles in a Service Definition	675
	Specifying QoS Parameter Instances in a Service Definition	675
	Creating a Parameter Instance in a Profile	676
	Specifying QoS Parameter Instances in a Service Definition	676
	Modifying QoS Configurations with Service Manager	677
	Modifying Parameter Instances	677
	Modifying QoS Configurations in a Single Service Manager Event	679
	Modifying QoS Configurations by Using Other Sources	679
	Removing QoS Configurations Referenced by Service Manager	681
	QoS for Service Manager Considerations	682
	RADIUS or Service Manager	682
	Interoperability with Other Service Components	682
	QoS Statistics	682
	Ranges	682
	Configuring the Service Manager License	683
	Managing and Activating Service Sessions Overview	684
	Managing Subscriber Service Sessions by Using RADIUS Overview	684
	Activating Subscriber Service Sessions Using RADIUS	685
	Understanding Service Manager RADIUS Attributes	686
	Using Tags with RADIUS Attributes	689
	Deactivating Service Sessions Using RADIUS	690
	Setting Time or Volume Thresholds for a Service	690
	Using the Deactivate-Service Attribute	691
	Activating and Deactivating Subscriber Services by Using Mutex Groups Overview	691

Activating and Deactivating Multiple Services	692
Configuring a Mutex Service	693
Combined and Independent IPv4 and IPv6 Services in a Dual Stack Overview . .	694
Activation and Deactivation of IPv4 and IPv6 Services in a Dual Stack	696
Independent IPv4 and IPv6 Services in a Dual Stack	696
Combined IPv4 and IPv6 Service in a Dual Stack	696
Performance Impact on the Router and Compatibility with Previous Releases for an IPv4 and IPv6 Dual Stack	697
Understanding RADIUS Accounting for Service Manager	697
Service Interim Accounting Overview	699
Configuring Service Interim Accounting	702
Specifying the Service Accounting Interval	702
Specifying the User Accounting Interval	702
Configuring Calculation of Service Session Accounting Based on Scheduler Profiles Instead of Rate-Limit Profiles in Hierarchical Parent Groups for Forwarded Packets	703
Service Interim Accounting for IPv4 and IPv6 Services in a Dual Stack Overview	705
Collecting IPv6 Service Accounting Statistics Separately for Dual-Stack Subscribers Overview	706
Managing Subscriber Service Sessions by Using the CLI Overview	707
Activating Subscriber Service Sessions by Using the CLI Overview	708
Activating Subscriber Sessions by Using the CLI	709
Activating a Service for an Existing Subscriber	709
Creating and Activating a Service for a Subscriber	710
Preprovisioning Service Sessions	711
Service Session Profiles Overview	711
Working with Service Session Profiles	712
Creating a New Service Session Profile	713
Specifying Statistics Collection Settings	713
Specifying the Maximum Bandwidth for a Service Session	714
Specifying the Interval for the Active State of a Service Session	714
Deactivating Subscriber Service Sessions by Using the CLI Overview	715
Gracefully Deactivating Subscriber Service Sessions	716
Gracefully Deactivating Service Sessions Based on Owner Details	716
Gracefully Deactivating Service Sessions Based on Subscriber Details . . .	716
Forcing Immediate Deactivation of Subscriber Service Sessions	717
Using Service Session Profiles to Deactivate Service Sessions	717
Configuring Service Manager Statistics	718
Setting Up the Service Definition File for Statistics Collection	719
Enabling Statistics Collection with RADIUS	720
Enabling Statistics Collection with the CLI	721
Setting Up the External Parent Group Statistics Collection	722
Service Manager Performance Considerations	723
Example: Tiered Service Definition	724
Example: Video-on-Demand Service Definition	726
Example: Voice-over-IP Service Definition	728
Guided Entrance Service Overview	730
Example: Guided Entrance Service Definition	731

	Using COA Messages with Guided Entrance Services	734
	Configuring the HTTP Local Server to Support Guided Entrance	735
	Configuring the HTTP Local Server to Support Guided Entrance for IPv4 Subscribers	736
	Configuring the HTTP Local Server to Support Guided Entrance for IPv6 Subscribers	737
	Redirection of Subscriber Sessions When HTTP Local Server is Disabled or Not Configured	739
	Example: Combined IPv4 and IPv6 Service in a Dual Stack Service Definition . .	740
	Preservation of the Original URL During Redirection of Subscriber Sessions . .	746
	Configuring the Preservation of the Original URL During Redirection of Subscriber Sessions	747
Chapter 29	Monitoring Service Manager	749
	Setting a Baseline for HTTP Local Server Statistics	749
	Monitoring the Connections to the HTTP Local Server	750
	Monitoring the Configuration of the HTTP Local Server	750
	Monitoring Statistics for Connections to the HTTP Local Server	751
	Monitoring Profiles for the HTTP Local Server	752
	Monitoring the Default Interval for Interim Accounting of Services	753
	Monitoring the Status of the Service Manager License	753
	Monitoring Profiles for Service Manager	754
	Monitoring IPv4 and IPv6 Interfaces for Service Manager	755
	Monitoring Service Definitions	765
	Monitoring Service Session Profiles	766
	Monitoring Active Owner Sessions with Service Manager	767
	Monitoring Active Subscriber Sessions with Service Manager	770
	Monitoring the Number of Active Subscriber and Service Sessions with Service Manager	773
	Verifying Computation of Service Session Accounting Based on Scheduler Profiles	773
Part 7	Index	
	Index	777

List of Figures

Part 1	Managing Remote Access	
Chapter 1	Remote Access Overview	3
	Figure 1: Local Address Pool Hierarchy	25
	Figure 2: Shared Local Address Pools	27
Chapter 2	Configuring Remote Access	65
	Figure 3: Single PPP Clients per ATM Subinterface	83
	Figure 4: Multiple PPP Clients per ATM Subinterface	84
Part 2	Managing RADIUS and TACACS+	
Chapter 5	Configuring RADIUS Dynamic-Request Server	227
	Figure 5: Sample Remote Access Network Using RADIUS	228
Chapter 6	Configuring RADIUS Relay Server	235
	Figure 6: RADIUS Relay Server	236
Part 3	Managing L2TP	
Chapter 12	L2TP Overview	321
	Figure 7: Using the E Series Router as an LAC	322
	Figure 8: Using the E Series Router as an LNS	322
Chapter 13	Configuring an L2TP LAC	331
	Figure 9: Lockout States	358
Chapter 15	Configuring L2TP Dial-Out	407
	Figure 10: Network Model for Dial-Out	408
Part 4	Managing DHCP	
Chapter 19	DHCP Local Server Overview	471
	Figure 11: Non-PPP Equal Access via the Router	474
Chapter 20	Configuring DHCP Local Server	485
	Figure 12: Non-PPP Equal-Access Configuration Example	507
Chapter 21	Configuring DHCP Relay	511
	Figure 13: Passing 802.1p Values to the DHCP Server	536
Chapter 22	Configuring the DHCP External Server Application	545
	Figure 14: DHCP External Server	546

Part 5	Managing the Subscriber Environment	
Chapter 24	Configuring Subscriber Management	605
	Figure 15: DHCP External Server	608
Chapter 26	Configuring Subscriber Interfaces	619
	Figure 16: Example of a Dynamic Interface Stack	621
	Figure 17: Example of a Dynamic Subscriber Interface	621
	Figure 18: Subscriber Interfaces over Ethernet	624
	Figure 19: Subscriber Interfaces in a Cable Modem Network	627
	Figure 20: Associating Subnets with a VPN Using Subscriber Interfaces	628
	Figure 21: IP over Ethernet Dynamic Subscriber Interface Configuration	630
	Figure 22: Subscriber Interfaces Using a Destination Address to Demultiplex Traffic	636
	Figure 23: Subscriber Interfaces Using a Source Address to Demultiplex Traffic	639
	Figure 24: IP over Ethernet Dynamic Subscriber Interface Configuration	643
	Figure 25: IP over VLAN over Ethernet Dynamic Subscriber Interface Configuration	645
	Figure 26: IP over Bridged Ethernet over ATM Dynamic Subscriber Interface Configuration	647
	Figure 27: GRE Tunnel Dynamic Subscriber Interface Configuration	649
Part 6	Managing Subscriber Services	
Chapter 28	Configuring Service Manager	661
	Figure 28: Service Manager Configuration Flowchart	667
	Figure 29: Sample Service Definition Macro File	670
	Figure 30: QoS Configuration Dependency Chain	681
	Figure 31: Comparing RADIUS Login and RADIUS COA Methods	685
	Figure 32: Guided Entrance	731
	Figure 33: Input Traffic Flow with Rate-Limit Profile on an External Parent Group for a Combined IPv4/IPv6 Service	741
	Figure 34: Output Traffic Flow with Rate-Limit Profile on an External Parent Group for a Combined IPv4/IPv6 Service	741

List of Tables

	About the Documentation	xxxiii
	Table 1: Notice Icons	xxxiv
	Table 2: Text and Syntax Conventions	xxxiv
Part 1	Managing Remote Access	
Chapter 1	Remote Access Overview	3
	Table 3: Username and Domain Name Examples	13
	Table 4: aaa strip-domain Example	16
	Table 5: Local UDP Port Ranges by RADIUS Request Type	19
	Table 6: RADIUS IETF Attributes in Preauthentication Request	33
	Table 7: VSAs That Apply to Dynamic IP Interfaces	34
	Table 8: Traffic-Shaping VSAs That Apply to Dynamic IP Interfaces	36
	Table 9: Supported RADIUS Acct-Terminate-Cause Codes	37
	Table 10: RADIUS Attributes Specifying LAG Interface	46
	Table 11: SRC Client and COPS Terminology	47
Chapter 3	Monitoring and Troubleshooting Remote Access	111
	Table 12: show aaa accounting Output Fields	115
	Table 13: show aaa accounting vr-group Output Fields	119
	Table 14: show aaa strip-domain Output Fields	120
	Table 15: show aaa domain-map Output Fields	121
	Table 16: show aaa profile Output Fields	124
	Table 17: show aaa route-download Output Fields	126
	Table 18: show aaa route-download routes Output Fields	128
	Table 19: show aaa route-download routes global Output Fields	130
	Table 20: show aaa statistics Output Fields	132
	Table 21: show configuration category aaa global-attributes Output Fields	135
	Table 22: show configuration category aaa local-authentication Output Fields	136
	Table 23: show configuration category aaa server-attributes include-defaults Output Fields	137
	Table 24: show cops info Output Fields	139
	Table 25: show cops statistics Output Fields	141
	Table 26: show ip local alias Output Fields	143
	Table 27: show ip local pool Output Fields	144
	Table 28: show ip local shared-pool Output Fields	145
	Table 29: show radius override Output Fields	147
	Table 30: show radius servers Output Fields	148
	Table 31: show radius statistics Output Fields	151
	Table 32: show ssc info Output Fields	156

	Table 33: show sssc statistics Output Fields	158
	Table 34: show sssc option Output Fields	160
	Table 35: show subscribers Output Fields	166
	Table 36: show terminate-code Output Fields	168
	Table 37: show ipv6 local pool Output Fields	169
	Table 38: show ipv6 local pool poolName Output Fields	170
	Table 39: show ipv6 local pool statistics Output Fields	171
	Table 40: show ipv6 local ndra-pool Output Fields	172
	Table 41: show ipv6 local ndra-pool poolName Output Fields	173
	Table 42: show ipv6 local ndra-pool statistics Output Fields	175
Part 2	Managing RADIUS and TACACS+	
Chapter 4	Configuring RADIUS Attributes	179
	Table 43: AAA Access Message RADIUS IETF Attributes Supported	183
	Table 44: AAA Access Message Juniper Networks (Vendor ID 4874) VSAs Supported	186
	Table 45: AAA Accounting Message RADIUS IETF Attributes Supported	194
	Table 46: AAA Accounting Message Juniper Network (Vendor ID 4874) VSAs Supported	197
	Table 47: AAA Accounting Tunnel Message RADIUS Attributes Supported	200
	Table 48: DSL Forum (Vendor ID 3561) VSAs Supported in AAA Access and Accounting Messages	206
	Table 49: CLI AAA Access Message RADIUS Attributes Supported	208
	Table 50: CLI Commands Used to Configure RADIUS IETF Attributes	210
	Table 51: CLI Commands Used to Configure Juniper Networks VSAs	213
	Table 52: ANCP (L2C)-Related Keywords for radius include Command	216
	Table 53: RADIUS Attributes Included in Corresponding RADIUS Messages	219
Chapter 5	Configuring RADIUS Dynamic-Request Server	227
	Table 54: Error-Cause Codes (RADIUS Attribute 101)	230
	Table 55: Error-Cause Codes (RADIUS Attribute 101)	233
Chapter 6	Configuring RADIUS Relay Server	235
	Table 56: Required RADIUS Access-Request Attributes	236
	Table 57: Required RADIUS Accounting Attributes	237
Chapter 7	RADIUS Attribute Descriptions	241
	Table 58: RADIUS IETF Attributes Supported by JunosE Software	241
	Table 59: Juniper Networks (Vendor ID 4874) VSA Formats	248
	Table 60: JunosE Software DSL Forum (Vendor ID 3561) VSA Formats	261
	Table 61: RADIUS Attribute Passed Through by JunosE Software	262
Chapter 8	Application Terminate Reasons	265
	Table 62: Default AAA Mappings	265
	Table 63: Default L2TP Mappings	267
	Table 64: Default PPP Mappings	283
	Table 65: Default RADIUS Client Mappings	290
Chapter 9	Monitoring RADIUS	291
	Table 66: show radius override Output Fields	292

	Table 67: show radius attributes-included Output Fields	297
	Table 68: show radius dynamic-request statistics Output Fields	299
	Table 69: show radius dynamic-request servers Output Fields	300
	Table 70: show radius relay statistics Output Fields	302
	Table 71: show radius relay servers Output Fields	303
	Table 72: show radius relay udp-checksum Output Fields	304
Chapter 10	Configuring TACACS+	305
	Table 73: TACACS-Related Terms	306
	Table 74: TACACS+ Accounting Information	308
Chapter 11	Monitoring TACACS+	315
	Table 75: show statistics tacacs Output Fields	316
	Table 76: show tacacs Output Fields	317
Part 3	Managing L2TP	
Chapter 12	L2TP Overview	321
	Table 77: L2TP Terms	322
Chapter 14	Configuring an L2TP LNS	365
	Table 78: L2TP-Resynch-Method RADIUS Attribute	389
	Table 79: Transmit Connect Speeds for L2TP over ATM 1483 Example	397
	Table 80: Transmit Connect Speeds for L2TP over Ethernet Example	398
	Table 81: Tunnel--Tx-Speed-Method RADIUS Attribute	402
Chapter 15	Configuring L2TP Dial-Out	407
	Table 82: L2TP Dial-Out Terms	408
	Table 83: Chassis Operational States	411
	Table 84: Virtual Router Operational States	411
	Table 85: Target Operational States	411
	Table 86: Session Operational States	412
	Table 87: Additions to RADIUS Attributes in Access-Accept Messages	414
Chapter 16	L2TP Disconnect Cause Codes	421
	Table 88: PPP Disconnect Cause Codes	421
Chapter 17	Monitoring L2TP and L2TP Dial-Out	427
	Table 89: show aaa domain-map Output Fields	428
	Table 90: show aaa tunnel-group Output Fields	430
	Table 91: show aaa tunnel-parameters Output Fields	432
	Table 92: show l2tp Output Fields	434
	Table 93: show l2tp destination Output Fields	436
	Table 94: show l2tp destination lockout Output Fields	437
	Table 95: show l2tp destination profile Output Fields	439
	Table 96: show l2tp destination summary Output Fields	441
	Table 97: show l2tp received-disconnect-cause-summary Output Fields	442
	Table 98: show l2tp session Output Fields	443
	Table 99: show l2tp session summary Output Fields	444
	Table 100: show l2tp switch-profile Output Fields	445
	Table 101: show l2tp tunnel Output Fields	446
	Table 102: show l2tp tunnel summary Output Fields	449

	Table 103: show l2tp dial-out Output Fields	451
	Table 104: show l2tp dial-out session Output Fields	455
	Table 105: show l2tp dial-out target Output Fields	456
	Table 106: show l2tp dial-out virtual-router Output Fields	457
Part 4	Managing DHCP	
Chapter 19	DHCP Local Server Overview	471
	Table 107: Local Pool Selection in Equal-Access Mode	473
	Table 108: Local Pool Selection in Standalone Mode Without AAA Authentication	475
	Table 109: Local Pool Selection in Standalone Mode with AAA Authentication	475
Chapter 21	Configuring DHCP Relay	511
	Table 110: Router Configuration and Transmission of DHCP Reply Packets	517
	Table 111: Effect of Commands on Option 82 Suboption Settings	527
Chapter 23	Monitoring and Troubleshooting DHCP	559
	Table 112: show ip dhcp-local excluded Output Fields	562
	Table 113: show dhcp binding Output Fields	564
	Table 114: show dhcp count Output Fields	567
	Table 115: show dhcp host Output Fields	569
	Table 116: show ip dhcp-external binding Output Fields	570
	Table 117: show ip dhcp-external binding-id	571
	Table 118: show ip dhcp-local binding Output Fields	572
	Table 119: show ip dhcp-external configuration Output Fields	573
	Table 120: show ip dhcp-external statistics Output Fields	574
	Table 121: show dhcp-external Output Fields	574
	Table 122: show ip dhcp-local pool Output Fields	576
	Table 123: show ip dhcp-local auth Output Fields	578
	Table 124: show ip dhcp-local Output Fields	579
	Table 125: show ip dhcp-local leases Output Fields	580
	Table 126: show ip dhcp-local statistics output fields	582
	Table 127: show dhcp vendor-option Output Fields	584
	Table 128: show ip dhcp-capture Output Fields	585
	Table 129: show dhcp relay Output Fields	586
	Table 130: show dhcp relay proxy statistics Output Fields	587
	Table 131: show dhcp relay statistics Output Fields	589
	Table 132: show dhcp server statistics Output Fields	591
	Table 133: show dhcp server Output Fields	592
	Table 134: show ipv6 dhcpv6-local binding Output Fields	593
	Table 135: show ipv6 dhcpv6-local dns-domain-searchlist Output Fields	594
	Table 136: show ipv6 dhcpv6-local dns-servers Output Fields	594
	Table 137: show ipv6 dhcpv6-local prefix-lifetime Output Fields	595
	Table 138: show ipv6 dhcpv6-local statistics Output Fields	596
	Table 139: show ipv6 dhcpv6-local auth config Output Fields	597
	Table 140: show ip dhcp-local limits Output Fields	598
	Table 141: show ip dhcp-local reserved Output Fields	599
	Table 142: show dhcp summary Output Fields	599

	Table 143: show dhcp proxy-client binding Output Fields	600
	Table 144: show ipv6 dhcpv6-local limits Output Fields	601
Part 5	Managing the Subscriber Environment	
Chapter 25	Monitoring Subscriber Management	615
	Table 145: show ip service-profile Output Fields	615
	Table 146: show ip-subscriber Output Fields	617
Chapter 27	Monitoring Subscriber Interfaces	655
	Table 147: show ip demux interface Output Fields	655
	Table 148: show ip-subscriber Output Fields	657
Part 6	Managing Subscriber Services	
Chapter 28	Configuring Service Manager	661
	Table 149: Service Manager Terms and Acronyms	663
	Table 150: JunosE Objects Tracked by Service Manager	669
	Table 151: Sample Modifications Using the Add and Initial-Value Keywords	678
	Table 152: Sample Modifications Using Parameter Instances	678
	Table 153: Configuration Within a Single Service Manager Event	679
	Table 154: Modifying QoS Configurations with Other Sources	680
	Table 155: Service Manager RADIUS Attributes	687
	Table 156: Sample RADIUS Access-Accept Packet	688
	Table 157: Using Tags	689
	Table 158: Service Manager RADIUS Accounting Attributes	698
	Table 159: Determining the Service Interim Accounting Interval	700
	Table 160: Sample Acct-Start Message for a Service Session	700
	Table 161: RADIUS-Enabled Statistics	720
	Table 162: Sample RADIUS Attributes	725
	Table 163: Sample RADIUS Attributes	727
	Table 164: Sample RADIUS Attributes	729
	Table 165: Sample RADIUS Attributes	734
	Table 166: Deactivating a Guided Entrance Service	735
Chapter 29	Monitoring Service Manager	749
	Table 167: show ip http scalar Output Fields	750
	Table 168: show ip http server Output Fields	751
	Table 169: show ip http statistics Output Fields	752
	Table 170: show profile Output Fields	753
	Table 171: show aaa service accounting interval Output Fields	753
	Table 172: show license service-management Output Fields	754
	Table 173: show profile Output Fields	754
	Table 174: show ip interface Output Fields	757
	Table 175: show ipv6 interface Output Fields	760
	Table 176: show service-management service-definition Output Fields	766
	Table 177: show service-management service-session-profile Output Fields	767
	Table 178: show service-management owner-session Output Fields	768
	Table 179: show service-management subscriber-session Output Fields	771
	Table 180: show service-management summary Output Fields	773

About the Documentation

- E Series and JunosE Documentation and Release Notes on page xxxiii
- Audience on page xxxiii
- E Series and JunosE Text and Syntax Conventions on page xxxiii
- Obtaining Documentation on page xxxv
- Documentation Feedback on page xxxv
- Requesting Technical Support on page xxxvi

E Series and JunosE Documentation and Release Notes

For a list of related JunosE documentation, see
<http://www.juniper.net/techpubs/software/index.html>.

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

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

Audience

This guide is intended for experienced system and network specialists working with Juniper Networks E Series Broadband Services Routers in an Internet access environment.

E Series and JunosE Text and Syntax Conventions

Table 1 on page xxxiv defines notice icons used in this documentation.

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 xxxiv defines text and syntax conventions that we use throughout the E Series and JunosE documentation.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents commands and keywords in text.	<ul style="list-style-type: none"> Issue the clock source command. Specify the keyword exp-msg.
Bold text like this	Represents text that the user must type.	host1(config)#traffic class low-loss1
Fixed-width text like this	Represents information as displayed on your terminal's screen.	host1#show ip ospf 2 Routing Process OSPF 2 with Router ID 5.5.0.250 Router is an Area Border Router (ABR)
<i>Italic text like this</i>	<ul style="list-style-type: none"> Emphasizes words. Identifies variables. Identifies chapter, appendix, and book names. 	<ul style="list-style-type: none"> There are two levels of access: <i>user</i> and <i>privileged</i>. <i>clusterId</i>, <i>ipAddress</i>. <i>Appendix A, System Specifications</i>
Plus sign (+) linking key names	Indicates that you must press two or more keys simultaneously.	Press Ctrl + b.

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

Convention	Description	Examples
Syntax Conventions in the Command Reference Guide		
Plain text like this	Represents keywords.	terminal length
<i>Italic text like this</i>	Represents variables.	<i>mask, accessListName</i>
(pipe symbol)	Represents a choice to select one keyword or variable to the left or to the right of this symbol. (The keyword or variable can be either optional or required.)	diagnostic line
[] (brackets)	Represent optional keywords or variables.	[internal external]
[]* (brackets and asterisk)	Represent optional keywords or variables that can be entered more than once.	[level1 level2 l1]*
{ } (braces)	Represent required keywords or variables.	{ permit deny } { in out } { clusterId ipAddress }

Obtaining Documentation

To obtain the most current version of all Juniper Networks technical documentation, see the Juniper Networks TechLibrary at <http://www.juniper.net/techpubs/>.

To download complete sets of technical documentation to create your own documentation CD-ROMs or DVD-ROMs, see the Portable Libraries page at

<http://www.juniper.net/techpubs/resources/index.html>

Copies of the Management Information Bases (MIBs) for a particular software release are available for download in the software image bundle from the Juniper Networks website at <http://www.juniper.net/>.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation to better meet your needs. Send your comments to techpubs-comments@juniper.net, or fill out the documentation feedback form at <https://www.juniper.net/cgi-bin/docbugreport/>. If you are using e-mail, be sure to include the following information with your comments:

- Document or topic name
- URL or page number
- Software release version

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>
- Download the latest versions of software and review release notes: <http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <http://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

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

Opening a Case with JTAC

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

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

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

PART 1

Managing Remote Access

- [Remote Access Overview on page 3](#)
- [Configuring Remote Access on page 65](#)
- [Monitoring and Troubleshooting Remote Access on page 111](#)

CHAPTER 1

Remote Access Overview

- [Remote Access Overview on page 4](#)
- [Remote Access Platform Considerations on page 6](#)
- [Remote Access References on page 6](#)
- [Mapping a User Domain to a Virtual Router Overview on page 7](#)
- [Domain Name and Realm Name Overview on page 9](#)
- [Example: Domain Name and Realm Name on page 13](#)
- [Example: Stripping Domain Name Per Virtual Router for RADIUS Server Authentication on page 14](#)
- [Single Name Specification for Users from a Domain Overview on page 16](#)
- [RADIUS Authentication and Accounting Servers Configuration Overview on page 17](#)
- [SNMP Traps and System Log Messages Overview on page 22](#)
- [Local Authentication Servers Configuration Overview on page 23](#)
- [Tunnel Subscriber Authentication Configuration Overview on page 23](#)
- [Name Server Addresses Configuration Overview on page 24](#)
- [Local Address Servers Configuration Overview on page 25](#)
- [DHCP Features on page 28](#)
- [Domain Name Aliases Overview on page 28](#)
- [AAA Profile Configuration Overview on page 29](#)
- [RADIUS Route-Download Server for Route Distribution Overview on page 29](#)
- [AAA Logical Line Identifier for Subscriber Tracking Overview on page 31](#)
- [RADIUS Attributes in Preauthentication Request on page 33](#)
- [Considerations for Using the LLID on page 33](#)
- [VSAs for Dynamic IP Interfaces Overview on page 34](#)
- [Mapping Application Terminate Reasons and RADIUS Terminate Codes Overview on page 36](#)
- [Timeout Configuration Overview on page 38](#)
- [Standard RADIUS IPv6 Attributes for IPv6 Neighbor Discovery Router Advertisements and DHCPv6 Prefix Delegation Configuration on page 39](#)

- [Maximum Number of IPv6 Prefixes Assigned to Clients by Using Only the DHCPv6 Local Server on page 40](#)
- [Maximum Number of IPv6 Prefixes Assigned to Clients by Using Both DHCPv6 Local Server and Neighbor Discovery Router Advertisements on page 40](#)
- [Duplicate IPv6 Prefix Check Overview on page 42](#)
- [Duplicate IPv6 Prefix Detection in the AAA User Profile Database Overview on page 42](#)
- [Guidelines for Duplicate Address Verification on page 43](#)
- [Propagation of LAG Subscriber Information to AAA and RADIUS on page 45](#)
- [SRC Client Configuration Overview on page 47](#)
- [SRC Client and COPS Terminology on page 47](#)
- [Retrieval of DSL Line Rate Information from Access Nodes Overview on page 50](#)
- [DHCPv6 Local Address Pools for Allocation of IPv6 Prefixes Overview on page 51](#)
- [Example: Delegating the DHCPv6 Prefix on page 54](#)
- [Sending of Delegated-IPv6-Prefix from the DHCPv6 Local Server to SRC Client Overview on page 56](#)
- [IPv6 Prefix Allocation Using Neighbor Discovery Router Advertisements from IPv6 Address Pools Overview on page 57](#)
- [Understanding IPCP and IPv6CP Negotiations for IPv4 and IPv6 Clients Based on RADIUS-Returned Attributes on page 60](#)
- [Sending the Acct-Start and Acct-Update Messages After NCP Negotiation Is Completed Overview on page 63](#)

Remote Access Overview

Broadband Remote Access Server (B-RAS) is an application running on your router that:

- Aggregates the output from digital subscriber line access multiplexers (DSLAMs)
- Provides user Point-to-Point Protocol (PPP) sessions or IP-over-Asynchronous Transfer Mode (ATM) sessions
- Enforces quality of service (QoS) policies
- Routes traffic into an Internet service provider's (ISP's) backbone network

A DSLAM collects data traffic from multiple subscribers into a centralized point so that it can be uploaded to the router over an ATM connection via a DS3, OC3, E3, or OC12 link.

The router provides the logical termination for PPP sessions, as well as the interface to authentication and accounting systems.

The following sections provide an overview of remote access:

- [B-RAS Data Flow on page 5](#)
- [Configuring IP Addresses for Remote Clients on page 5](#)
- [AAA Overview on page 5](#)

B-RAS Data Flow

The router performs several tasks for a digital subscriber line (DSL) PPP user to establish a PPP connection. This is an example of the way B-RAS data might flow:

1. Authenticate the subscriber using RADIUS authentication.
2. Assign an IP address to the PPP/IP session via RADIUS, local address pools, or Dynamic Host Configuration Protocol (DHCP).
3. Terminate the PPP encapsulation or tunnel a PPP session.
4. Provide user accounting via RADIUS.



NOTE: For information about configuring RADIUS attributes see the *Configuring RADIUS Attributes* chapter.

Configuring IP Addresses for Remote Clients

A remote client can obtain an IP address from one of the following:

- RADIUS server
- Local address server
- DHCP proxy client and server
- DHCP relay agent (Bridged IP only)
- DHCP local server
- DHCP external server

For information about configuring DHCP support on the E Series router, see the *DHCP Overview* chapter.

For information about how to configure a RADIUS server, see your RADIUS server documentation.

AAA Overview

Collectively, authentication, authorization, and accounting are referred to as AAA. Each has an important but separate function.

- **Authentication**—Determines who the user is, then determines whether that user should be granted access to the network. The primary purpose is to prevent intruders from networks. It uses a database of users and passwords.
- **Authorization**—Determines what the user is allowed to do by giving network managers the ability to limit network services to different users.
- **Accounting**—Tracks what the user did and when they did it. You can use accounting for an audit trail or for billing for connection time or resources used.

Central management of AAA means the information is in a single, centralized, secure database, which is much easier to administer than information distributed across numerous devices.

Related Documentation

- [Remote Access Configuration Tasks on page 66](#)

Remote Access Platform Considerations

B-RAS services are supported on all E Series routers.

For information about the modules supported on E Series routers:

- See the *ERX Module Guide* for modules supported on ERX7xx models, ERX14xx models, and the ERX310 Broadband Services Router.
- See the *E120 and E320 Module Guide* for modules supported on the Juniper Networks E120 and E320 Broadband Services Routers.
- [B-RAS Protocol Support on page 6](#)

B-RAS Protocol Support

The E Series router supports the following protocols for B-RAS services:

- PPP
- PPP over Ethernet (PPPoE)
- Bridged Ethernet
- Layer 2 Tunneling Protocol (L2TP), both L2TP access concentrator (LAC) and L2TP network server (LNS)

Related Documentation

- [Remote Access Overview on page 4](#)
- [Remote Access References on page 6](#)

Remote Access References

For more information about the topics covered in this chapter, see the following documents:

- RFC 2748—The COPS (Common Open Policy Service) Protocol (January 2000)
- RFC 2865—Remote Authentication Dial In User Service (RADIUS) (June 2000)
- RFC 3084—COPS Usage for Policy Provisioning (COPS-PR) (March 2001)
- RFC 3159—Structure of Policy Provisioning Information (SPPI) (August 2001)
- RFC 3198—Terminology for Policy-Based Management (November 2001)

- RFC 3317—Differentiated Services Quality of Service Policy Information Base (DIFFSERV-PIB)
- RFC 3318—Framework Policy Information Base (March 2003)

JunosE Release Notes, Appendix A, System Maximums—Refer to the Release Notes corresponding to your software release for information about the number of concurrent RADIUS requests that the router supports for authentication and accounting servers.

**Related
Documentation**

- [Remote Access Overview on page 4](#)
- [Remote Access Platform Considerations on page 6](#)

Mapping a User Domain to a Virtual Router Overview

You can configure RADIUS authentication, accounting, and local address pools for a specific virtual router and then map a user domain to that virtual router.

The router keeps track of the mapping between domain names and virtual-routers. Use the **aaa domain-map** command to map a user domain to a virtual router.



NOTE: This domain name is not the NT domain sometimes found on the Dialup Networking dialog box.

When the router is configured to require authentication of a PPP user, the router checks for the appropriate user domain-name-to-virtual-router mapping. If it finds a match, the router sends a RADIUS authentication request to the RADIUS server configured for the specific virtual router.

The following sections describe how to map a user domain to a virtual router:

- [Mapping User Requests Without a Valid Domain Name on page 7](#)
- [Mapping User Requests Without a Configured Domain Name on page 8](#)
- [Using DNIS on page 8](#)
- [Redirected Authentication on page 8](#)
- [IP Hinting on page 9](#)

Mapping User Requests Without a Valid Domain Name

You can create a mapping between a domain name called **default** and a specific virtual router so that the router can map usernames that contain a domain name that does not have an explicit map.

If a user request is submitted with a domain name for which the router cannot find a match, the router looks for a mapping between the domain name **default** and a virtual router. If a match is found, the user's request is processed according to the RADIUS server configured for the named virtual router. If no entry is found that maps **default** to a specific virtual router, the router sends the request to the RADIUS server configured on the default virtual router.

Mapping User Requests Without a Configured Domain Name

You can map a domain name called **none** to a specific virtual router so that the router can map usernames that do not contain a domain name.

If a user request is submitted without a domain name, the router looks for a mapping between the domain name **none** and a virtual router. If a match is found, the user's request is processed according to the RADIUS server configured for the named virtual router. If the router does not find the domain name **none**, it checks for the domain name **default**. If no matching entries are found, the router sends the request to the server configured on the default virtual router.

Using DNIS

The E Series router supports dialed number identification service (DNIS). With DNIS, if users have a called number associated with them, the router searches the domain map for the called number. If it finds a match, the router uses the matching domain map entry information to authenticate the user. If the router does not find a match, it searches the domain map using normal processing.



NOTE: For DNIS to work, the router must be acting as the LNS. Also, the phone number configured in the **aaa domain-map** command must be an exact match to the value passed by L2TP in the called number AVP (AVP 21).

For example, as specified in the following sequence, a user calling 9785551212 would be terminated in **vrouter_88**, while a user calling 8005554433 is terminated in **vrouter_100**.

```
host1(config)#aaa domain-map 9785551212 vrouter_88
host1(config)#aaa domain-map 8005554433 vrouter_100
```

Redirected Authentication

Redirected authentication provides a way to offload AAA activity on the router, by providing the domain-mapping-like feature remotely on the RADIUS server. Redirected authentication works as follows:

1. The router sends an authentication request (in the form of a RADIUS access-request message) to the RADIUS server that is configured in the default VR.
2. The RADIUS server determines the user's AAA VR context and returns this information in a RADIUS response message to the router.
3. The router then behaves in similar fashion as if it had received the VR context from the local domain map.

To maintain local control, the only VR allowed to redirect authentication is the default VR. Also, to prevent loopbacks, the redirection may occur only once to a non-default VR.

To maintain flexibility, the redirection response may include idle time or session attributes that are considered as default unless the redirected authentication server overrides them.

For example, if the RADIUS server returns the VR context along with an idle timeout attribute with the value set to 20 minutes, the router uses this idle timeout value unless the RADIUS server configured in the VR context returns a different value.

Since the router supports the RADIUS User-Name attribute [1] in the RADIUS response message, the default VR RADIUS server may override the user's name (this can be a stripped name or an entirely different name). Overriding is useful for the case when the user enters a login name containing a domain name that is significant only to the RADIUS server in the default VR.

IP Hinting

You can allocate an address before authentication of PPP sessions. This address is included in the Access-Request sent to the authentication server as an IP address hint.

Related Documentation

- [Domain Name and Realm Name Overview on page 9](#)

Domain Name and Realm Name Overview

To provide flexibility in how the router handles different types of usernames, the software lets you specify the part of a username to use as the domain name, how the domain name is designated, and how the router parses names. It also allows you to set whether or not the router strips the domain name from the username before it sends the username to the RADIUS server.

By default, the router parses usernames as follows:

realmName/personalName@domainName

The string to the left of the forward slash (/) is the realm name, and the string to the right of the at-symbol (@) is the domain name. For example, in the username juniper/jill@abc.com, juniper is the realm name and abc.com is the domain name.

The router allows you to:

- Use the realm name as the domain name.
- Use delimiters other than / to designate the realm name.
- Use delimiters other than @ to designate the domain name.
- Use either the domain or the realm as the domain name when the username contains both a realm and domain name.
- Change the direction in which the router searches for the domain name or the realm name.

To provide these features, the router allows you to specify delimiters for the domain name and realm name. You can use up to eight one-character delimiters each for domain and realm names. The router also lets you specify how it parses usernames to determine which part of a username to use as the domain name.

The following sections describe domain name and realm name:

- [Using the Realm Name as the Domain Name on page 10](#)
- [Using Delimiters Other Than @ on page 10](#)
- [Using Either the Domain or the Realm as the Domain Name on page 10](#)
- [Specifying the Domain Name or Realm Name Parse Direction on page 11](#)
- [Stripping the Domain Name on page 11](#)
- [Stripping the Domain Name Per Virtual Router on page 11](#)

Using the Realm Name as the Domain Name

Typically, a realm appears before the user field and is separated with the / character; for example, usEast/jill@abc.com. To use the realm name usEast rather than abc.com as the domain name, set the realm name delimiter to /. For example:

```
host1(config)#aaa delimiter realmName /
```

This command causes the router to use the string to the left of the / as the domain name. If the realm name delimiter is null (the default), the router will not search for the realm name.

Using Delimiters Other Than @

You can set up the router to recognize delimiters other than @ to designate the domain name. Suppose there are two users: bob@abc.com and pete!xyz.com, and you want to use both of their domain names. In this case you would set the domain name delimiter to @ and !. For example:

```
host1(config)#aaa delimiter domainName @!
```

Using Either the Domain or the Realm as the Domain Name

If the username contains both a realm name and a domain name delimiter, you can use either the domain name or the realm name as the domain name. As previously mentioned, the router treats usernames with multiple delimiters as though the realm name is to the left of the realm delimiter and the domain name is to the right of the domain delimiter.

If you set the parse order to:

- domain-first—The router searches for a domain name first. For example, for username usEast/lori@abc.com, the domain name is abc.com.
- realm-first—The router searches for a realm name first and uses the realm name as the user's domain name. For username usEast/lori@abc.com, the domain is usEast.

For example, if you set the delimiter for the realm name to / and set the delimiter for the domain name to @, the router parses the realm first by default. The username usEast/lori@abc.com results in a domain name of usEast. To cause the parsing to return abc.com as the domain, enter the **aaa parse-order domain-first** command.

Specifying the Domain Name or Realm Name Parse Direction

You can specify the direction—either left to right or right to left—in which the router performs the parsing operation when identifying the realm name or domain name. This feature is particularly useful if the username contains nested realm or domain names. For example, for a username of `userjohn@abc.com@xyz.com`, you can identify the domain as either `abc.com@xyz.com` or as `xyz.com`, depending on the parse direction that you specify.

You use either the **left-to-right** or **right-to-left** keywords with one of the following keywords to specify the type of search and parsing that the router performs:

- **domainName**—The router searches for the next domain delimiter value in the direction specified. When it reaches a delimiter, the router uses anything to the right of the delimiter as the domain name. Domain parsing is from right to left by default.
- **realmName**—The router searches for the next realm delimiter value in the direction specified. When it reaches a delimiter, the router uses anything to the left of the delimiter as the realm name. Realm parsing is from left to right by default.
- Example

```
host1(config)#aaa parse-direction domainName left-to-right
```

Stripping the Domain Name

The router provides a feature that strips the domain name from the username before it sends the username to the RADIUS server in an Access-Request message. You can enable or disable this feature by using the **strip-domain** command.

The router does not permanently remove the domain name from the username of the subscriber. Instead, the router blocks the domain name from being sent in the Access-Request message, thereby retaining the domain name. So the output of the **show subscribers** command always displays the complete username, including the domain name (if available), of the subscriber regardless of the status of the strip domain feature.

By default, the domain name is the text after the last @ character. However, if you changed the domain name parsing using the **aaa delimiter**, **aaa parse-order**, or **aaa parse direction** commands, the router strips the domain name and delimiter that result from the parsing.

Stripping the Domain Name Per Virtual Router

The **aaa domain-map** command maps a domain name to a virtual router. It determines the authentication and accounting access for all subscribers belonging to a particular domain. However, if a subscriber profile is configured for a virtual router using the **ppp authentication** command, the authentication for the virtual router configured at the profile level takes priority over the one configured at the domain level. If multiple profiles from the same domain are being used, the subscribers may end up in different virtual routers for authentication.

In such a scenario, you can use the **aaa strip-domain** command to strip a part of the username of the subscriber. The resulting username is then used as the new username for that subscriber for RADIUS authentication and accounting.

The router does not permanently remove the domain name from the username of the subscriber. Instead, the router blocks the domain name from being sent in the Access-Request message, thereby retaining the domain name. So the output of the **show subscribers** command always displays the complete username, including the domain name (if available), of the subscriber regardless of the status of the strip domain feature.



NOTE:

- The **aaa strip-domain** command can be configured on nondefault virtual routers only.
- When you execute the **show subscribers** and **logout subscribers** commands with the **username** keyword, you must specify the complete username, including the domain name (if available), of the subscriber regardless of the status of the strip domain feature on a virtual router or AAA domain map.

Subscriber Username for RID, COA Requests, and Lawful Intercepts When Strip Domain Is Enabled

When strip domain is enabled for a virtual router, the username used to identify the subscriber session for RADIUS Initiated Disconnect (RID), Change of Authorization (COA), and lawful intercepts requests is the same as the subscriber username sent to the RADIUS server for authentication.

For example, if a subscriber with the username `user1@123.com$test1` has a resulting username of `user1@123.com` due to the strip domain configuration, then the username for all the incoming RID and COA requests and the lawful intercept requests is `user1@123.com`.

Using the Strip Domain Functionality Per Virtual Router When Strip Domain Is Enabled for an AAA Domain Map

When strip domain is enabled for an AAA domain map using the **strip-domain enable** command in Domain Map Configuration mode, the strip domain configured for a virtual router may cause the username stripping to happen twice depending on the configuration.

For example, consider a subscriber with the username `user1@test.com$test1$test2`. Consider the following configurations for a domain map:

```
host1(config)#aaa domain-map test2
host1(config-domain-map)#strip-domain enable
```

The following has also been configured on the nondefault virtual router:

```
host1(config)#aaa strip-domain enable
host1(config)#aaa strip-domain delimiter domainname $
```

In this example, when the domain name is stripped for the subscriber with the username `user1@test.com$test1$test2`, the resulting string that is sent for RADIUS authentication is `user1`. Thus, when strip domain is configured for a domain map as well as a nondefault virtual router, depending on the configurations, the domain name may be stripped twice: once at the virtual router level and then at the domain map level.

To prevent the domain name from being stripped twice for the same subscriber, you must ensure that the strip domain functionality is configured appropriately for the domain map and for the nondefault virtual router.

Redirected Authentication When Strip Domain Is Enabled

Strip domain configured on a virtual router does not work in case of a redirected authentication. In a redirected authentication, the RADIUS server sends an Access-Accept message for a subscriber from the virtual router on which the subscriber is already authenticated.

For example, on virtual router `vr1`, the `aaa strip-domain` is configured. A subscriber with the username `user1@123.com` is already authenticated on `vr1` using the RADIUS server authentication. Now, if you send an Access-Request message to authenticate the same subscriber on `vr1`, the Access-Request message carries the original username, `user1@123.com`, and renders strip domain ineffective during redirected authentication.

Related Documentation

- [Example: Domain Name and Realm Name on page 13](#)
- [Example: Stripping Domain Name Per Virtual Router for RADIUS Server Authentication on page 14](#)

Example: Domain Name and Realm Name

This section provides examples of possible domain or realm name results that you might obtain, depending on the commands and options you specify. This example uses the following username:

username: usEast/userjohn@abc.com@xyz.com

The router is configured with the following commands:

```
host1(config)#aaa delimiter domainName @!
host1(config)#aaa delimiter realmName /
```

[Table 3 on page 13](#) shows the username and domain name that result from the parsing action of the various commands.

Table 3: Username and Domain Name Examples

Command	Resulting Username	Resulting Domain Name
<code>aaa parse-order realm-first</code>	<code>userjohn@abc.com@xyz.com</code>	<code>usEast</code>
<code>aaa parse-order domain-first</code>	<code>userjohn@abc.com</code>	<code>xyz.com</code>

Table 3: Username and Domain Name Examples (*continued*)

Command	Resulting Username	Resulting Domain Name
<code>aaa parse-direction domainName right-to-left</code>	userjohn@abc.com	xyz.com
<code>aaa parse-direction domainName left-to-right</code>	userjohn	abc.com@xyz.com
<code>aaa parse-direction realmName right-to-left</code>	userjohn@abc.com@xyz.com	usEast
<code>aaa parse-direction realmName left-to-right</code>	userjohn@abc.com@xyz.com	usEast

Related Documentation • [Domain Name and Realm Name Overview on page 9](#)

Example: Stripping Domain Name Per Virtual Router for RADIUS Server Authentication

This example illustrates the final username for a subscriber, based on the virtual router applied.

1. Configure the five virtual routers.

```
host(config)#profile VR1
host(config-profile)#ppp authentication virtual-router vr1 pap chap
host(config-profile)#exit
host(config)#profile VR2
host(config-profile)#ppp authentication virtual-router vr2 pap chap
host(config-profile)#exit
host(config)#profile VR3
host(config-profile)#ppp authentication virtual-router vr3 pap chap
host(config-profile)#exit
host(config)#profile VR4
host(config-profile)#ppp authentication virtual-router vr4 pap chap
host(config-profile)#exit
host(config)#profile VR5
host(config-profile)#ppp authentication virtual-router vr2 pap chap
host(config-profile)#exit
```

2. Access the context of a previously created virtual router and enable the strip domain functionality for each virtual router.

```
host(config)#virtual-router vr1
host:vr1(config)#aaa strip-domain enable
host:vr1(config)#aaa strip-domain delimiter domainName $
host:vr1(config)#aaa strip-domain parse-direction domainName left-to-right
host:vr1(config)#radius authentication server 10.209.154.193
host:vr1(config)#key bras
host:vr1(config)#exit
host:vr1(config)#radius accounting server 10.209.154.193
host:vr1(config-radius)#key bras
```

```
host:vr1(config-radius)#exit
host:vr1(config)#virtual-router vr2
```

```
host:vr2(config)#aaa strip-domain enable
host:vr2(config)#aaa strip-domain parse-direction domainName left-to-right
host:vr2(config)#radius authentication server 10.209.154.194
host:vr2(config-radius)#key bras
host:vr2(config-radius)#exit
host:vr2(config)#radius accounting server 10.209.154.194
host:vr2(config-radius)#key bras
host:vr2(config-radius)#exit
host:vr2(config)#virtual-router vr3
```

```
host:vr3(config)#radius authentication server 10.209.154.193
host:vr3(config-radius)#key bras
host:vr3(config-radius)#exit
host:vr3(config)#radius accounting server 10.209.154.193
host:vr3(config-radius)#key bras
host:vr3(config-radius)#exit
host:vr3(config)#virtual-router vr4
```

```
host:vr4(config)#aaa strip-domain enable
host:vr4(config)#aaa strip-domain delimiter domainName %
host:vr4(config)#radius authentication server 10.209.154.194
host:vr4(config-radius)#key bras
host:vr4(config-radius)#exit
host:vr4(config)#radius accounting server 10.209.154.195
host:vr4(config-radius)#key bras
host:vr4(config-radius)#exit
host:vr4(config)#virtual-router vr5
```

```
host:vr5(config)#aaa strip-domain enable
host:vr5(config)#radius authentication server 10.209.154.193
host:vr5(config-radius)#key bras
host:vr5(config-radius)#exit
host:vr5(config)#radius accounting server 10.209.154.192
host:vr5(config-radius)#key bras
host:vr5(config-radius)#exit
```

Based on the configurations of the virtual routers, [Table 4 on page 16](#) lists the final username that is sent to the RADIUS server for RADIUS authentication and accounting for each virtual router.

**NOTE:**

- The output of the *show subscribers* command does not display the final username of the subscriber. It displays the complete username, including the domain name (if available), of the subscriber for all virtual routers regardless of the status of the strip domain feature.
- When you execute the *show subscribers* and *logout subscribers* commands with the *username* keyword, you must specify the complete username, including the domain name (if available), of the subscriber regardless of the status of the strip domain feature on a virtual router or AAA domain map.

Table 4: aaa strip-domain Example

Subscribers	Virtual Router Applied	Final Username
user1@123.com\$test	vr1	user1@123.com
user2@123.com\$test	vr2	user2
user3@123.com\$test	vr3	user3@123.com\$test
user4@123.com%test	vr4	user4@123.com
user5@123.com@test\$test	vr5	user5@123.com

Related Documentation

- [Mapping a User Domain to a Virtual Router Overview on page 7](#)
- [Domain Name and Realm Name Overview on page 9](#)

Single Name Specification for Users from a Domain Overview

Assigning a single username and a single password for all users associated with a domain provides better compatibility with some RADIUS servers. You can use this feature for domains that require the router to tunnel, but not terminate, PPP sessions.

When users request a PPP session, they specify usernames and passwords. During the negotiations for the PPP session, the router authenticates legitimate users.



NOTE: This feature works only for users authenticated by Password Authentication Protocol (PAP) and not by Challenge Handshake Authentication Protocol (CHAP).

If you configure this feature, the router substitutes the specified username and password for all authenticated usernames and passwords associated with that domain.

There are two options for this feature. The router can:

- Substitute the domain name for each username and one new password for each existing password.

For example, if the domain name is xyz.com and you specify the password xyz_domain, the router associates the username xyz.com and the password xyz_domain with all users from xyz.com.

- Substitute one new username for each username and one new password for each existing password.

For example, if the domain name is xyz.com and you specify the username xyz_group and the password xyz_domain, the router associates these identifiers with all users from xyz.com.

To use a single username and a single password for all users from a domain:

1. Access Domain Map Configuration mode using the **aaa domain-map** command.
2. Specify the new username and password using the **override-user** command.

**Related
Documentation**

- [Example: Associating all Subscribers of a PPP Interface with a Specific Domain Name on page 86](#)

RADIUS Authentication and Accounting Servers Configuration Overview

Initially, a RADIUS client sends a request to a RADIUS authentication or accounting server. The RADIUS server uses the configured IP address, the UDP port number, and the secret key to make the connection. The RADIUS client waits for a response for a configurable timeout period and then retransmits the request. The RADIUS client retransmits the request for a user-configurable retry limit.

- If there is no response from the primary RADIUS server, the RADIUS client submits the request to the secondary RADIUS server using the timeout period and retry limit configured for the secondary RADIUS server.
- If the connection attempt fails for the secondary RADIUS server, the router submits the request to the tertiary server and so on until it either is granted access on behalf of the client or there are no more configured servers.
- If another authentication server is not configured, the router attempts the next method in the method list; for accounting server requests, the information is dropped.

For example, suppose that you have configured the following authentication servers: Auth1, Auth2, Auth3, Auth4, and Auth5. Your router attempts to send an authentication request to Auth1. If Auth1 is unavailable, the router submits the request to Auth2, then Auth3, and so on until an available server is found. If Auth5, the last configured authentication server, is not available, the router attempts the next method in the methods list. If the only method configured is RADIUS, then the router notifies the client that the request has been denied.

**NOTE:**

- The number of RADIUS servers you can configure depends on available memory.
- The order in which you configure servers determines the order in which the router contacts those servers on behalf of clients.

The following sections explain how to configure RADIUS authentication and accounting servers:

- [Server Access on page 18](#)
- [Server Request Processing Limit on page 18](#)
- [Authentication and Accounting Methods on page 19](#)
- [Supporting Exchange of Extensible Authentication Protocol Messages on page 20](#)
- [Immediate Accounting Updates on page 21](#)
- [Interim Accounting Updates on page 21](#)
- [Duplicate and Broadcast Accounting on page 21](#)

Server Access

The router offers two options by which servers are accessed:

- **Direct**—The first authentication or accounting server that you configure is treated as the primary authentication or accounting server, the next server configured is the secondary, and so on.
- **Round-robin**—The router sends the initial request to the first configured authentication or accounting server, the next request to the second configured server, and so on until the last configured server. After sending the request to the last configured server, the router again starts this cycle from the first configured server.

Use the **radius algorithm** command to specify the server access method.

When you configure the first RADIUS accounting server, a RADIUS Acct-On message is sent. When you delete the last accounting server, a RADIUS Acct-Off message is sent.

Server Request Processing Limit

You can configure RADIUS authentication servers and accounting servers to use different UDP ports on the router. This enables the same IP address to be used for both an authentication server and an accounting server. However, you cannot use the same IP address for multiple authentication servers or for multiple accounting servers.



NOTE: For information about the number of concurrent RADIUS requests that the router supports for authentication and accounting servers, see *JunosE Release Notes, Appendix A, System Maximums*.

The E Series router listens to a range of UDP source (or local) ports for RADIUS responses. Each UDP source port supports a maximum of 255 RADIUS requests. When the 255 per-port limit is reached, the router opens the next source port. When the **max-sessions** command limit is reached, the router submits the request to the next configured server.

[Table 5 on page 19](#) lists the range of UDP ports the router uses for each type of RADIUS request.

Table 5: Local UDP Port Ranges by RADIUS Request Type

RADIUS Request Type	ERX310, ERX710, ERX1410, and E120 Broadband Services Routers	ERX1440 and E320 Broadband Services Routers
RADIUS authentication	50000–50124	50000–50124
RADIUS accounting	50125–50249	50125–50499
RADIUS preauthentication	50250–50374	50500–50624
RADIUS route-download	50375–50500	50625–50749

Authentication and Accounting Methods

When you configure authentication, authorization, and accounting (AAA) services for your B-RAS environment, one important task is to specify the authentication and accounting method used. The JunosE Software gives you the flexibility to configure authentication or accounting methods based on the type of subscriber. This feature allows you to enable RADIUS authentication for some subscribers, while disabling authentication completely for other subscribers. Similarly, you can enable RADIUS accounting for some subscribers, but no accounting for others. For example, you might use RADIUS authentication for ATM 1483 subscribers, while granting IP subscriber management interfaces access without authentication (using the **none** keyword).

You can specify the authentication or accounting method you want to use, or you can specify multiple methods in the order in which you want them used. For example, if you specify the **radius** keyword followed by the **none** keyword when configuring authentication, AAA initially attempts to use RADIUS authentication. If no RADIUS servers are available, AAA uses no authentication. The JunosE Software currently supports **radius** and **none** as accounting methods and **radius**, **none**, and **local** as authentication methods. See [“Local Authentication Servers Configuration Overview” on page 23](#) for information about local authentication.

You can configure authentication and accounting methods based on the following types of subscribers:

- ATM 1483
- Tunnels (for example, L2TP tunnels)
- Point-to-Point Protocol (PPP)

- RADIUS relay server
- IP subscriber management interfaces



NOTE: IP subscriber management interfaces are static or dynamic interfaces that are created or managed by the JunosE Software's subscriber management feature.

Supporting Exchange of Extensible Authentication Protocol Messages

Extensible Authentication Protocol (EAP) is a protocol that supports multiple methods for authenticating a peer before allowing network layer protocols to transmit over the link. JunosE Software supports the exchange of EAP messages between JunosE applications, such as PPP, and an external RADIUS authentication server.

The JunosE Software's AAA service accepts and passes EAP messages between the JunosE application and the router's internal RADIUS authentication server. The internal RADIUS authentication server, which is a RADIUS client, provides EAP pass-through—the RADIUS client accepts the EAP messages from AAA, and sends the messages to the external RADIUS server for authentication. The RADIUS client then passes the response from the external RADIUS authentication server back to the AAA service, which then sends a response to the JunosE application. The AAA service and the internal RADIUS authentication service do not process EAP information—both simply act as pass-through devices for the EAP message.

The router's local authentication server and TACACS+ authentication servers do not support the exchange of EAP messages. These type of servers deny access if they receive an authentication request from AAA that includes an EAP message. EAP messages do not affect the **none** authentication configuration, which always grants access.

The local RADIUS authentication server uses the following RADIUS attributes when exchanging EAP messages with the external RADIUS authentication server:

- Framed-MTU (attribute 12)—Used if AAA passes an MTU value to the internal RADIUS client
- State (attribute 24)—Used in Challenge-Response messages from the external server and returned to the external server on the subsequent Access-Request
- Session-Timeout (attribute 27)—Used in Challenge-Response messages from the external server
- EAP-Message (attribute 79)—Used to fragment EAP strings into 253-byte fragments (the RADIUS limit)
- Message-Authenticator (attribute 80)—Used to authenticate messages that include an EAP-Message attribute

For additional information on configuring PPP to use EAP authentication, see *JunosE Link Layer Configuration Guide*.

Immediate Accounting Updates

You can use the **aaa accounting immediate-update** command to configure immediate accounting updates on a per-VR basis. If you enable this feature, the E Series router sends an Acct-Update message to the accounting server immediately on receipt of a response (ACK or timeout) to the Acct-Start message.

This feature is disabled by default. Use the **enable** keyword to enable immediate updates and the **disable** keyword to halt them.

The accounting update contains 0 (zero) values for the input/output octets/packets and 0 (zero) for uptime. If you have enabled duplicate or broadcast accounting, the accounting update goes to both the primary virtual router context and the duplicate or broadcast virtual router context.

Interim Accounting Updates

You can use the **aaa accounting interim-update** command to enable or disable the interim user accounting updates feature on a per-virtual router basis. If you enable this feature, AAA periodically sends an Interim-Acct request at a configured user accounting interval to a primary accounting server. You can use the **aaa user accounting interval** command to configure the user accounting interval. When the user accounting interval is not configured (by default, the interval is set to zero), AAA does not send the Interim-Acct request to the primary accounting server even if the interim user accounting updates feature is enabled.

The interim user accounting updates feature is enabled by default. You can use the **disable** keyword with the **aaa accounting interim-update** command to disable the interim user accounting updates feature, which disables sending of the Interim-Acct request even though the user accounting interval is configured. You can use the **enable** keyword with the **aaa accounting interim-update** command or the **no** version of the **aaa accounting interim-update** command to enable the interim accounting updates feature.

Duplicate and Broadcast Accounting

Normally, the JunosE Software sends subscriber-related AAA accounting information to the virtual router that authenticates the subscriber. If an operational virtual router is configured that is different from the authentication router, it also receives the accounting information. You can optionally configure duplicate or broadcast AAA accounting, which sends the accounting information to additional virtual routers simultaneously. The accounting information is always sent to the authenticating virtual router. The accounting information is sent to the operational virtual router only if duplicate accounting is not enabled and if authenticating virtual router is different than the operational virtual router.

Both the duplicate and broadcast accounting features are supported on a per-virtual router context, and enable you to specify particular accounting servers that you want to receive the accounting information.

For example, you might use broadcast accounting to send accounting information to a group of your private accounting servers. Or you might use duplicate accounting to send the accounting information to a customer's accounting server.

- Duplicate accounting—Sends the accounting information to a particular virtual router
- Broadcast accounting—Sends the accounting information to a group of virtual routers. An accounting virtual router group can contain up to four virtual routers and the E Series router supports a maximum of 100 virtual router groups. The accounting information continues to be sent to the duplicate accounting virtual router, if one is configured.

UDP Checksums

Each virtual router on which you configure B-RAS is enabled to perform UDP checksums by default. You can disable and reenable UDP checksums.

Related Documentation

- [Remote Access Configuration Tasks on page 66](#)
- [Configuring Interim, Broadcast, and Policy-Based Accounting in Virtual Router Groups and PPP Profiles on page 89](#)

SNMP Traps and System Log Messages Overview

The router can send Simple Network Management Protocol (SNMP) traps to alert network managers when:

- A RADIUS server fails to respond to a request.
- A RADIUS server that previously failed to respond to a request (and was consequently removed from the list of active servers) returns to active service.

Returning to active service means that the E Series RADIUS client receives a valid response to an outstanding RADIUS request after the server is marked unavailable.

- All RADIUS servers within a VR context fail to respond to a request.

The router also generates system log messages when RADIUS servers fail to respond or when they return to active service; no configuration is required for system log messages.

The following sections describe SNMP Traps and system log messages:

- [SNMP Traps on page 22](#)
- [System Log Messages on page 23](#)

SNMP Traps

The router generates SNMP traps and system log messages as follows:

- If the first RADIUS server fails to respond to the RADIUS request, the E Series RADIUS client issues a system log message and, if configured, an SNMP trap indicating that the RADIUS server timed out. The E Series RADIUS client will not issue another system log message or SNMP trap regarding this RADIUS server until the dead time expires, if configured, or for 3 minutes if dead time is not configured.
- The E Series RADIUS client then sends the RADIUS request to the second configured RADIUS server. If the second RADIUS server fails to respond to the RADIUS request,

the E Series RADIUS client again issues a system log message and, if configured, an SNMP trap indicating that the RADIUS server timed out.

- This process continues until either the E Series RADIUS client receives a valid response from a RADIUS server or the list of configured RADIUS servers is exhausted. If the list of RADIUS servers is exhausted, the E Series RADIUS client issues a system log message and, if configured, an SNMP trap indicating that all RADIUS servers have timed out.

If the E Series RADIUS client receives a RADIUS response from a “dead” RADIUS server during the dead time period, the RADIUS server is restored to active status.

If the router receives a valid RADIUS response to an outstanding RADIUS request, the E Series client issues a system log message and, if configured, an SNMP trap indicating that the RADIUS server is now available.

System Log Messages

You do not need to configure system log messages. The router automatically sends them when individual servers do not respond to RADIUS requests and when all servers on a VR fail to respond to requests. The following are the formats of the warning level system log messages:

```
RADIUS [ authentication | accounting ] server serverAddress unavailable in VR
virtualRouterName [; trying nextServerAddress]
RADIUS no [ authentication | accounting ] servers responding in VR virtualRouterName
RADIUS [ authentication | accounting ] server serverAddress available in VR
virtualRouterName
```

- Related Documentation**
- [Configuring SNMP Traps on page 71](#)

Local Authentication Servers Configuration Overview

The AAA local authentication server enables the E Series router to provide local PAP and CHAP user authentication for subscribers. The router also provides limited authorization, using the IP address, IP address pool, and operational virtual router parameters. When a subscriber logs on to the E Series router that is using local authentication, the subscriber is authenticated against user entries in a local user database; the optional parameters are assigned to subscribers after the subscriber is authenticated.

- Related Documentation**
- [Creating the AAA Local Authentication Environment on page 73](#)
 - [Creating AAA Local User Databases on page 73](#)

Tunnel Subscriber Authentication Configuration Overview

When a AAA domain map includes any tunnel configuration, users in this domain are considered to be tunnel subscribers. By default, any such subscriber is granted access without being authenticated by the authentication server. Access is granted even when the user provides an invalid username and password. The tunnel configuration for the subscriber comes from the AAA domain map.

For example, if the authentication protocol for a AAA domain map is RADIUS, AAA grants access to subscribers from this domain immediately without sending access requests to the configured RADIUS server. Because of this behavior, these subscribers cannot get any additional control attributes from the authentication server. This reduces your ability to manage the tunnel subscribers.

In this default situation, if you want the domain subscribers to be managed by the authentication server for any control attribute, then that domain map cannot have any tunnel configuration. Typically, this means you must configure the subscriber individually.

You can use the **tunnel-subscriber authentication** command to get around this limitation. When you enable authentication with this command, access requests for the tunnel subscribers in the domain are sent to the configured authentication server. When the access replies from authentication server are processed, various user attributes from the server can be applied to the subscribers.

When the authentication server returns tunnel attributes, these returned values take precedence over the corresponding local tunnel configuration values in the AAA domain map. If the server does not return any tunnel attributes, then the tunnel subscriber's tunnel settings are configured according to the domain map's tunnel settings.

If the authentication server returns a redirect VSA and the corresponding AAA domain map has local tunnel configurations, the VSA is ignored. Access is denied to the user when the authentication server rejects the access request.

The **tunnel-subscriber authentication** command has no effect on subscribers in a domain with no tunnel configuration. When a AAA domain map has no tunnel configuration, subscribers in the domain are authenticated by the authentication server. If the server grants access, then the subscribers get their tunnel settings only from the authentication server.

By default, tunnel subscribers in the domain are granted access with no external authentication. Use the **enable** keyword to enable authentication. Use the **disable** keyword to restore disable user authentication.

To configure authentication of tunnel subscribers within a AAA domain by an external authentication server.

- Example

```
host1(config-domain-map)#tunnel-subscriber authentication enable
```

**Related
Documentation**

- [Mapping a User Domain to a Virtual Router Overview on page 7](#)
- *tunnel-subscriber authentication*

Name Server Addresses Configuration Overview

You can assign IP or IPv6 addresses for DNS and IP addresses for WINS name servers. During setup negotiations between the router and remote PC clients using PPP (Internet Protocol Control Protocol [IPCP] specifically), the remote client may request the DNS

and WINS server IP addresses. If the IP addresses passed to the router by the remote PC client are different from the ones configured on your router, the router returns the values that you configured as the correct values to the remote PC client. This behavior is controlled by the **ppp peer dns** and **ppp peer wins** interface commands.

If a PPP client request contains address values of 0.0.0.0 for the name servers, the router considers that the remote PC client is not configured and returns the configured values as the correct values to the remote PC client.

The DNS and WINS addresses are considered as part of the PPP user information. These addresses are provided to the PPP client as part of the IPCP negotiations between PPP peers. For details, see RFC 1877—PPP Internet Protocol Control Protocol Extensions for Name Server Addresses (December 1995).



NOTE: All name server address parameters are defined in the context of a virtual router.

Related Documentation

- *ppp peer*


Local Address Servers Configuration Overview

The local address server allocates IP addresses from a pool of addresses stored locally on the router. You can optionally configure shared local address pools to obtain addresses from a DHCP local address pool that is in the same virtual router. Addresses are provided automatically to client sessions requiring an IP address from a virtual router that is configured to use a local address pool.

A local address server is defined in the context of a virtual router. You create a local address server when you configure the first local pool. Local address servers exist as long as the virtual router exists or until you remove them by deleting all configured pools.

Figure 1 on page 25 illustrates the local address pool hierarchy. Multiple local address server instances, one per virtual router, can exist. Each local address server can have one or more local address pools. Each pool can contain a number of IP addresses that are available for allocation and used by clients, such as PPP sessions.

Figure 1: Local Address Pool Hierarchy



The following sections describe local address servers:

- [Local Address Pool Ranges on page 26](#)
- [Local Address Pool Aliases on page 26](#)
- [Shared Local Address Pools on page 26](#)
- [SNMP Thresholds on page 27](#)

Local Address Pool Ranges

As shown in [Figure 1 on page 25](#), each local address pool is named and contains ranges of sequentially ordered IP addresses. These addresses are allocated when the AAA server makes a request for an IP address.

If a local address pool range is exhausted, the next range of addresses is used. If all pool ranges are exhausted, you can configure a new range to extend or supplement the existing range of addresses, or you can create a new pool. The newly created pool range is then used for future address allocation. If addresses allocated from the first pool range are released, then subsequent requests for addresses are taken from the first pool range.

Addresses are assigned sequentially from a range within a pool. If a range has no addresses available, the next range within that pool is used. If a pool has no addresses available, the next configured pool is used, unless a specific pool is indicated.



NOTE: The AAA server does not allocate Class D and Class E addresses to the subscribers from the local pool, even if they are included in the local address pool range.

Local Address Pool Aliases

An alias is an alternate name for an existing local address pool. It comprises an alias name and a pool name.

When the AAA server requests an IP address from a specific local address pool, the local address server first verifies whether an alias exists for the requested pool. If an alias exists, the IP address is allocated from the pool specified by the alias. If no alias exists, the IP address is allocated from the pool originally specified in the request.

The use of aliases simplifies management of subscribers. For example, you can use an alias to migrate subscribers from one local address pool to another. Instead of having to modify countless subscriber records on the AAA server, you create an alias to make the configuration change.

Shared Local Address Pools


Typically, the local address server allocates IP addresses from a pool of addresses that is stored locally on the router. However, *shared* local address pools enable a local address server to hand out addresses that are allocated from DHCP local server address pools within the same virtual router. The addresses are configured and managed within DHCP.

Therefore, thresholds are not configured on the shared pool, but are instead managed by the referenced DHCP local server pool.

A shared local address pool references one DHCP address pool. The shared local address pool can then obtain addresses from the referenced DHCP address pool and from any DHCP address pools that are linked to the referenced DHCP address pool.

Figure 2 on page 27 illustrates a shared local address pool environment that includes four linked DHCP address pools. In the figure, both Shared_LAS_Pool_A and Shared_LAS_Pool_B reference DHCP_Pool_1, and can therefore obtain addresses from all four DHCP address pools. Shared_LAS_Pool_C references DHCP_Pool_3 and can get addresses from DHCP_Pool_3 and DHCP_Pool_4.

Figure 2: Shared Local Address Pools



When the local address server requests an address from a shared address pool, the address is returned from the referenced DHCP pool or a subsequent linked pool. If no address is available, DHCP notifies the local address server and the search is ended.

Keep the following guidelines in mind when using shared local address pools:

- The DHCP attributes do not apply to shared local address pools; for example, the lease time for shared local address pools is infinite.
- When you delete the referenced DHCP address pool, DHCP notifies the local address server and logs out all subscribers that are using addresses from the deleted pool.
- When you delete a shared local address pool, the local address server logs out the subscribers that are using addresses from the deleted pool, then notifies DHCP and releases the addresses.
- If the chain of linked DHCP address pools is broken, no action is taken and the existing subscribers retain their address. However, the DHCP local address pools that are no longer part of the chain are now unable to provide any new addresses.

The following commands create the shared address pools in Figure 2 on page 27:

```

host1(config)#ip local shared-pool Shared_LAS_Pool_A DHCP_Pool_1
host1(config)#ip local shared-pool Shared_LAS_Pool_B DHCP_Pool_1
host1(config)#ip local shared-pool Shared_LAS_Pool_C DHCP_Pool_3
  
```

SNMP Thresholds

A local address pool has SNMP thresholds associated with it that enable the local address server to signal SNMP traps when certain conditions exist. These thresholds include high utilization threshold and abated utilization threshold. If the outstanding addresses of a

pool or a pool group exceed the high utilization threshold and the SNMP trap signaling is enabled, SNMP is notified. Likewise, when a pool's utilization drops below the abated utilization threshold, SNMP is notified.

A local address pool can be linked to a second local address pool so that when the first pool utilization reaches 100%, the DHCP local server uses the second pool. For generation of SNMP traps, the utilization of addresses is calculated for all the pools that are in the linked pools and they are collectively considered as an aggregated pool group.

**Related
Documentation**

- [Configuring a Local Address Server on page 81](#)

DHCP Features

DHCP provides a mechanism through which computers using Transmission Control Protocol/IP (TCP/IP) can obtain an IP address and protocol configuration parameters automatically from a DHCP server on the network.

The E Series router provides support for the following DHCP features:

- DHCP proxy client
- DHCP relay agent
- DHCP relay proxy
- DHCP local server
- DHCP external server

**Related
Documentation**

- [DHCP Overview Information on page 461](#)

Domain Name Aliases Overview

You can translate an original domain name to a new domain name via the **translate** command. The command allows you to create domain name aliases; that is, the grouping of multiple domain names into a single domain name. You can partition PPP subscribers with the same domain into separate domains, based on the PPP interface.



NOTE: Partitioning subscribers does not cause modification of a user's name or domain.

When you use aliases, you greatly simplify the configuration process. When there are a large number of domains and you use aliases, it reduces the configuration volume, thus requiring less NVS and memory usage.

**Related
Documentation**

- [Example: Associating all Subscribers of a PPP Interface with a Specific Domain Name on page 86](#)
- [Example: Associating Multiple Domain Names with a Specific Domain Name on page 87](#)

- *translate*

AAA Profile Configuration Overview

An AAA profile is a set of characteristics that act as a pattern that you can assign to domain names. Once you create an AAA profile, you can map it between a PPP client's domain name and certain AAA services on given interfaces. Using AAA profiles, you can:

- Allow or deny a domain name access to AAA authentication
- Map the original domain name to the mapped domain name for domain name lookup
- Use domain name aliases
- Force tunneling whenever a domain map contains tunnel attributes
- Manually set the NAS-Port-Type attribute (RADIUS attribute 61) for ATM and Ethernet interfaces
- Set the Service-Description attribute (RADIUS attribute 26-53)

An AAA profile contains a set of commands to control access for the incoming PPP subscriber. If no AAA profile is used, AAA continues as normal. The user's name and domain name are not changed as a result of an AAA profile mapping.



NOTE: There are two domain names with special meaning. The domain name **none** indicates that there is no domain name present in the subscriber's name. The domain name **default** indicates that no other match occurs.

Related Documentation

- [Single Name Specification for Users from a Domain Overview on page 16](#)
- [Example: Configuring AAA Local Authentication on page 77](#)

RADIUS Route-Download Server for Route Distribution Overview

The JunosE RADIUS route-download server provides periodic automatic distribution of IPv4 and IPv6 access routes, which enables preconfiguration and preadvertising of access routes before they are assigned to clients. Using the route-download server helps eliminate routing protocol storms and other delays in client service activation that can be caused by protocol convergence or a large number of simultaneous customer activations.

The RADIUS route-download server periodically sends a RADIUS Access-Request message to the RADIUS server to request that routes be downloaded. The RADIUS server then responds with an Access-Accept message and downloads the configured routes. When the download operation is complete, the route-download server installs the access routes in the routing table.

JunosE Software supports the creation of one RADIUS route-download server per chassis.

- [Format of Downloaded Routes on page 30](#)
- [How the Route-Download Server Downloads Routes on page 31](#)

Format of Downloaded Routes

The RADIUS server sends the downloaded routes to the RADIUS route-download server in the following format:

```
[ { vir | virtual-router } virtualRouterName ] [ vrf vrfName ] prefix-mask [ { null0 | null 0 } [ cost ] ] [ tag tagValue ]
```

For IPv4 routes, the route-download server accepts downloaded routes in either the Framed-Route attribute (RADIUS attribute 22) or the Cisco AV-pair attribute (Cisco VSA 26-1).

For IPv6 routes, the route-download server accepts downloaded routes in either the Framed-IPv6-Route attribute (RADIUS attribute 99) or the Cisco AV-pair attribute (Cisco VSA 26-1).

Framed-Route (RADIUS attribute 22)

```
NAS-1 Password = "14raddlsvr" User-Service-Type = Outbound-User
Framed-Route = "192.168.3.0 255.255.255.0 null0"
Framed-Route = "vrf vrfboston 192.168.1.0/24 null 0 0 tag 6"
Framed-Route = "vir host1 vrf vrfsunny 192.168.0.0/16 null0 0 tag 8"
```

Framed-IPv6-Route (RADIUS attribute 99)

```
NAS-1 Password = "14raddlsvr" User-Service-Type = Outbound-User
Framed-IPv6-Route = "2001:DB8:cc00:1::/48 null0"
Framed-IPv6-Route = "vrf test 2001:DB8:cc00:1::/48 null 0 0 tag 6"
Framed-IPv6-Route = "vir zzz vrf test1 2001:DB8:cc00:1::/48 null0 0 tag 8"
```

Cisco AV-Pair (Cisco VSA 26-1)

- NAS-1 Password = "14raddlsvr" User-Service-Type = Outbound-User
cisco-avpair = "ip:route = 192.168.3.0 255.255.255.0 null0"
cisco-avpair = "ip:route = vrf vrfboston 192.168.1.0/24 null 0 0 tag 6"
cisco-avpair = "ip:route = vir host1 vrf vrfsunny 192.168.0.0/16 null0 0 tag 8"
- NAS-1 Password = "14raddlsvr" User-Service-Type = Outbound-User
cisco-avpair = "ipv6:route=2001:DB8:cc00:1::/48 null0"
cisco-avpair = "ipv6:route=vrf test 2001:DB8:cc00:1::/48 null 0 0 tag 6"
cisco-avpair = "ipv6:route=vir zzz vrf test1 2001:DB8:cc00:1::/48 null0 0 tag 8"



NOTE: The prefix-mask entry in downloaded routes can be in the form of prefix length, prefix mask, or prefix. If prefix is used, the mask is determined by the IP address class of the prefix.

How the Route-Download Server Downloads Routes

The route-download server starts the initial route-download operation (for example, after a system reboot or the first time the route-download server is enabled) as soon as IP is established in the virtual router in which the download is performed. After the initial route-download process is established, the router repeats the route download operation based on either the default download schedule or the schedule you specify. You can also initiate an immediate route download at any time.

The RADIUS route-download server downloads routes in two stages—first, all routes are downloaded from the RADIUS server to the router's download database and examined for errors. Next, the router updates the routing table with the new routes, using the following guidelines:

- Adds all downloaded routes that are not already installed in the routing table
- Does not add downloaded routes that are already installed in the routing table
- Deletes routes from the routing table that do not appear in the newly downloaded group

Related Documentation

- [Configuring RADIUS AAA Servers on page 69](#)
- [Configuring the Route-Download Server to Download Routes on page 92](#)

AAA Logical Line Identifier for Subscriber Tracking Overview

You can configure the router to support the AAA logical line identification feature. This feature enables service providers to track subscribers on the basis of a virtual port known as the logical line ID (LLID).

The LLID is an alphanumeric string that logically identifies a subscriber line. The service provider maps each subscriber to an LLID based on the username and circuit ID from which the customer's calls originate. When a subscriber moves to a new physical line, the service provider's customer profile database is updated to map to the same LLID.

Because a subscriber's LLID remains the same regardless of the subscriber's physical location, using the LLID gives service providers a more secure mechanism for tracking subscribers and maintaining the customer database.

The following section explains how the router obtains and uses the LLID:

- [How the Router Obtains and Uses the LLID on page 31](#)

How the Router Obtains and Uses the LLID

To obtain an LLID for a subscriber, the router must issue two RADIUS access requests: a preauthentication request to obtain the LLID, followed by an authentication request encoded with the LLID returned in response to the preauthentication request.

To configure this feature, you:

1. Create an AAA profile that supports preauthentication (by using the **pre-authenticate** command in AAA Profile Configuration mode).
2. Specify the IP address of a RADIUS preauthentication server (by using the **radius pre-authentication server** command in Global Configuration mode) and of an authentication server (by using the **radius authentication server** command in Global Configuration mode).

The following steps describe how the router uses RADIUS to obtain and use the LLID. It is assumed that you have already configured an AAA profile for preauthentication and have defined both a RADIUS preauthentication server and a RADIUS authentication server. Typically, the preauthentication server and the authentication server reside in the same virtual router context in which the PPP subscriber is authenticated.

The router obtains and uses the LLID as follows:

1. A PPP subscriber requests authentication through RADIUS.
2. The router sends an Access-Request message to the RADIUS preauthentication server to obtain an LLID for the subscriber.

This step is referred to as the preauthentication request because it occurs before user authentication and authorization.

3. The preauthentication server returns the LLID to the router in the Calling-Station-Id (RADIUS attribute 31) of an Access-Accept message.

The router ignores any RADIUS attributes other than the Calling-Station-Id that are returned in the preauthentication Access-Accept message.

4. The router encodes the LLID in the RADIUS Calling-Station-Id and sends an Access-Request message to the RADIUS authentication server.

This step is referred to as the authentication request.

5. The RADIUS authentication server returns an Access-Accept message to the router that includes the tunnel attributes for the subscriber session.
6. For tunneled PPP subscribers, the router, acting as an L2TP access concentrator (LAC), encodes the LLID into L2TP Calling Number AVP 22 and sends this to the L2TP network server (LNS) in an incoming-call request (ICRQ) packet.

After a successful preauthentication request, the router always encodes the LLID in Calling Number AVP 22. The use of **aaa** commands such as **aaa tunnel calling-number-format** to control or change the inclusion of the LLID in Calling Number AVP 22 has no effect.

**Related
Documentation**

- [Configuring RADIUS AAA Servers on page 69](#)
- [Configuring the Router to Obtain the LLID for a Subscriber on page 93](#)

RADIUS Attributes in Preauthentication Request

Table 6 on page 33 describes the RADIUS IETF attributes that are always included in a preauthentication request to obtain the LLID. The attributes are listed in ascending order by standard number.

Table 6: RADIUS IETF Attributes in Preauthentication Request

Attribute Number	Attribute Name	Description
[1]	User-Name	Name of the user associated with the LLID, in the format: NAS-Port:<NAS-IP-Address>:<Nas-Port-Id> For example, nas-port:172.28.30.117:atm 4/1.104:2.104
[2]	User-Password	Password of the user to be authenticated; always set to “juniper”
[4]	NAS-IP-Address	IP address of the network access server (NAS) that is requesting authentication of the user; for example, 172.28.30.117
[5]	NAS-Port	Physical port number of the NAS that is authenticating the user; this is always interpreted as a bit field
[6]	Service-Type	Type of service the user has requested or the type of service to be provided; for example, framed
[61]	NAS-Port-Type	Type of physical port the NAS is using to authenticate the user
[77]	Connect-Info	Actual username; for example, jdoe@xyzcorp.east.com
[87]	NAS-Port-Id	Text string that identifies the physical interface of the NAS that is authenticating the user; for example, atm 4/1.104:2.104

The use of **radius** commands such as **radius calling-station-format** or **radius override calling-station-id** to control or change the inclusion of these attributes in the preauthentication request has no effect.

Related Documentation

- [RADIUS IETF Attributes on page 241](#)
- [Troubleshooting Subscriber Preauthentication on page 95](#)

Considerations for Using the LLID

The following considerations apply when you configure the router for subscriber preauthentication:

- Only PPP subscribers authenticating through RADIUS can use the AAA LLID feature on the router. PPP subscribers tunneled through domain maps cannot take advantage of this feature.
- The Calling-Station-Id [31] attribute is typically sent in RADIUS Access-Request messages, not in Access-Accept messages as is the case for this feature. As a result, your RADIUS server might require special configuration procedures to enable the Calling-Station-Id attribute to be returned in Access-Accept messages. See the documentation that came with your RADIUS server for information.
- The router ignores any RADIUS attributes other than the Calling-Station-Id that are returned in the preauthentication Access-Accept message.
- If a preauthentication request fails due to misconfiguration of the preauthentication server, timeout of the preauthentication server, or rejection of the preauthentication request by the preauthentication server, the authentication process continues normally and the preauthentication request is ignored.
- The router preserves the LLID value for established subscribers after a stateful SRP switchover.
- The **radius rollover-on-reject enable** command has no effect for a RADIUS preauthentication server. That is, you cannot use the **radius rollover-on-reject enable** command to configure the router to roll over to the next RADIUS preauthentication server when the router receives an Access-Reject message for the user it is authenticating.

**Related
Documentation**

- [Configuring RADIUS AAA Servers on page 69](#)

VSAs for Dynamic IP Interfaces Overview

Table 7 on page 34 describes the VSAs that apply to dynamic IP interfaces and are supported on a per-user basis from RADIUS. For details, see *JunosE Link Layer Configuration Guide*.

Table 7: VSAs That Apply to Dynamic IP Interfaces

VSA	Description	Type	Length	Subtype	Subtype Length	Value
Ingress-Policy-Name	Specifies the name of the input (ingress) policy	26	len	10	sublen	string: <i>input-policy-name</i>
Egress-Policy-Name	Specifies the name of the output (egress) policy	26	len	11	sublen	string: <i>output-policy-name</i>

Table 7: VSAs That Apply to Dynamic IP Interfaces (*continued*)

VSA	Description	Type	Length	Subtype	Subtype Length	Value
Ingress-Statistics	Indicates whether statistics are collected on input	26	12	12	6	integer: 0 – disable, 1 – enable
Egress-Statistics	Indicates whether statistics are collected on output	26	12	13	6	integer: 0 – disable, 1 – enable
QoS-Profile-Name	Specifies the name of the QoS profile to attach to the interface	26	len	26	sublen	string: <i>qos-profile-name</i>

To use the VSAs shown in [Table 7 on page 34](#):

- Specify the policy, or one or more QoS VSAs in the desired RADIUS user entries.
- Create the ingress or egress policy, or the QoS profile. Policies minimally consist of one or more policy commands and may include classifier control lists and rate limit profiles. See the *JunosE Policy Management Configuration Guide* for more information about policies and policy routing. See the *JunosE Quality of Service Configuration Guide* for information about creating QoS profiles.

When a dynamic interface is created according to a profile, the router checks with RADIUS to determine whether an input or output policy or a QoS profile must be applied to the interface. The VSA, if present, provides the name, enabling policy or QoS profile lookup. If found, the policy or QoS profile is applied to the dynamic interface.

The router also determines whether the creation profile specifies any policies to be applied to the interface. Policies specified by the RADIUS VSA supersede any specified by the profile, as described in the following example:

The RADIUS user entry includes an Ingress-Policy-Name VSA that specifies the policy input5. The profile specifies two policies, input7 and output1. In this case, the RADIUS-specified input policy (input5) and the profile-specified output policy (output1) are applied to the dynamic interface.

For information about assigning policies via profiles, see the *JunosE Policy Management Configuration Guide*. Only attributes assigned by RADIUS appear in RADIUS Acct-Start messages. RADIUS attributes specified by a profile for dynamic interfaces do not appear in RADIUS Acct-Start messages because the profile is not active when the Acct-Start message is generated. These attributes appear in RADIUS Acct-Stop messages for a profile that is active when the session is terminated.

The following section explains traffic shaping for PPP over ATM interfaces:

- [Traffic Shaping for PPP over ATM Interfaces on page 36](#)

Traffic Shaping for PPP over ATM Interfaces

The router supports the configuration of traffic shaping parameters for PPP over ATM (PPPoA) via domain-based profiles and RADIUS. In connection with this feature, [Table 8 on page 36](#) describes VSAs that apply to dynamic IP interfaces and are supported on a per-user basis from RADIUS.

Table 8: Traffic-Shaping VSAs That Apply to Dynamic IP Interfaces

VSA	Description	Type	Length	Subtype	Subtype Length	Value
Service-Category	Specifies the type of service	26	12	14	6	integer: 1 – UBR 2 – UBR PCR 3 – NRT VBR 4 – CBR 5 – RT VBR
PCR	Specifies the value for the peak cell rate (PCR)	26	12	15	6	integer
SCR	Specifies the value for the sustained cell rate (SCR)	26	12	16	6	integer
MBS	Specifies the maximum burst size (MBS)	26	12	17	6	integer

To configure traffic-shaping parameters for PPPoA via domain maps, use the **atm** command in Domain Map Configuration mode.

Related Documentation

- [Creating an IP Interface on page 83](#)

Mapping Application Terminate Reasons and RADIUS Terminate Codes Overview

The JunosE Software uses a default configuration that maps terminate reasons to RADIUS Acct-Terminate-Cause attributes. You can optionally create customized mappings between a terminate reason and a RADIUS Acct-Terminate-Cause attribute—these mappings enable you to provide different information about the cause of a termination.

When a subscriber's L2TP or PPP session is terminated, the router logs a message for the internal terminate reason and logs another message for the RADIUS Acct-Terminate-Cause attribute (RADIUS attribute 49). RADIUS attribute 49 is also

included in RADIUS Acct-Off and Acct-Stop messages. You can use the logged information to help monitor and troubleshoot terminated sessions.

Use the **show terminate-code** command to display information about the mappings between application terminate reasons and RADIUS Acct-Terminate-Cause attributes.

[Table 9 on page 37](#) lists the IETF RADIUS Acct-Terminate-Cause codes that you can use to map application terminate reasons. In addition, you can also configure and use proprietary codes for values beyond 22.

Table 9: Supported RADIUS Acct-Terminate-Cause Codes

Code	Name	Description
1	User Request	User initiated the disconnect (log out)
2	Lost Carrier	DCD was dropped on the port
3	Lost Service	Service can no longer be provided; for example, the user's connection to a host was interrupted
4	Idle Timeout	Idle timer expired
5	Session Timeout	Subscriber reached the maximum continuous time allowed for the service or session
6	Admin Reset	System administrator reset the port or session
7	Admin Reboot	System administrator terminated the session on the NAS; for example, prior to rebooting the NAS
8	Port Error	NAS detected an error on the port that required ending the session
9	NAS Error	NAS detected an error (other than on the port) that required ending the session
10	NAS Request	NAS ended the session for a non-error reason
11	NAS Reboot	NAS ended the session due to a non-administrative reboot
12	Port Unneeded	NAS ended the session because the resource usage fell below the low threshold; for example, the bandwidth-on-demand algorithm determined that the port was no longer needed
13	Port Preempted	NAS ended the session to allocate the port to a higher-priority use
14	Port Suspended	NAS ended the session to suspend a virtual session
15	Service Unavailable	NAS was unable to provide the requested service
16	Callback	NAS is terminating the current session in order to perform callback for a new session

Table 9: Supported RADIUS Acct-Terminate-Cause Codes (*continued*)

Code	Name	Description
17	User Error	An error in the user input caused the session to be terminated
18	Host Request	The login host terminated the session normally
19	Supplicant Restart	Supplicant state machine was reinitialized
20	Reauthentication Failure	A previously authenticated supplicant failed to reauthenticate successfully following expiration of the reauthentication timer or explicit reauthentication request by management action
21	Port Reinitialized	The port's MAC has been reinitialized
22	Port Administratively Disabled	The port has been administratively disabled

Related Documentation

- [Configuring Custom Mappings for PPP Terminate Reasons on page 95](#)

Timeout Configuration Overview

You can configure an idle timeout or a session timeout. The values you set are the default values for Point-to-Point Protocol Broadband Remote Access Server users. Attributes returned by RADIUS override these default settings on a per-user basis.

When you set an idle timeout, the PPP application on the router monitors both ingress (inbound) traffic and egress (outbound) traffic by default for the configured idle timeout period to determine whether to disconnect an inactive PPP session. If there is no activity in either direction on the interfaces for more than the configured idle timeout period, the router terminates the PPP session.

You can optionally configure the router to monitor only ingress traffic for the configured idle timeout period to determine session inactivity and subsequent disconnection of an inactive PPP session. Monitoring only ingress traffic for the idle timeout is useful for networks in which the PPP keepalive timer is disabled for wireless subscribers. Without the keepalive timer, the router cannot detect whether a wireless subscriber has been disconnected. Monitoring egress traffic does not indicate inactivity for wireless subscribers because egress traffic is always flowing. Enabling the router to monitor only ingress traffic enables you to selectively disconnect subscribers, including wireless subscribers, if no traffic is received for the configured idle timeout period.

If you do not configure a session timeout, or you set its value to 0, the session remains active for an infinite lifetime. You can use the **show ppp session-To-Thirteen-Years** command along with **show ppp interface full** in Privileged Exec or User Exec mode to verify whether the capability to preserve PPP sessions for a timeout duration of 13 years is enabled. If the **show ppp session-To-Thirteen-Years** command is not executed, the session timeout value is set to the maximum session timeout value of 366 days.

If the RADIUS server returns the value 0 for the Session-Timeout attribute, then the session remains active for an infinite lifetime even if a value is configured through the CLI.

The following sections describe timeout configuration:

- [Limiting Active Subscribers on page 39](#)
- [AAA Failure Notification for RADIUS on page 39](#)
- [Configuring AAA Session Timeout on page 39](#)

Limiting Active Subscribers

You can limit the number of active subscribers on a port or virtual router.

AAA Failure Notification for RADIUS

If a user passes RADIUS authentication, but fails AAA authentication, the RADIUS server may still allocate an address for the user from its internal address pool. To indicate to the RADIUS server to free the address, you can set up the router to send an Acct-Stop message if a user fails AAA.

Configuring AAA Session Timeout

You can use the **aaa timeout session *sessionTimeout*** command to configure a session timeout. Restoring the session timeout to the default value causes the PPP B-RAS session to remain active for an infinite lifetime.

Related Documentation

- [Configuring RADIUS AAA Servers on page 69](#)
- [Configuring Custom Mappings for PPP Terminate Reasons on page 95](#)

Standard RADIUS IPv6 Attributes for IPv6 Neighbor Discovery Router Advertisements and DHCPv6 Prefix Delegation Configuration

When an E Series router is configured for IP version 6, it uses router advertisements to announce its presence to other nodes connected to it. Hosts discover the addresses of their neighboring routers by listening for these advertisements. When the routing protocol process first starts on the server router, the server sends router advertisement packets every few seconds. Then, the server sends these packets less frequently. The server responds to route solicitation packets it receives from a client. The response is sent unicast, unless a router advertisement packet is due to be sent out momentarily. IPv6 supports the following router advertisement mechanisms:

- ICMPv6 Neighbor Discovery router advertisements
- DHCPv6 Prefix Delegation
- ICMPv6 Neighbor Discovery router advertisements followed by DHCPv6 Prefix Delegation

The AAA service on the router stores the prefixes that it receives from the RADIUS server during the PPPv6 authentication phase. After the PPPv6 link is established between the

subscriber and the B-RAS application running on the router, the router receives the ICMPv6 router solicitation message, the DHCPv6 Solicit message, or both of them based on the prefix advertisement mechanism. In previous releases, you were not able to configure the RADIUS attribute or VSA to be used for IPv6 Neighbor Discovery router advertisements and DHCPv6 Prefix Delegation through the CLI. As a result, the IPv6-NdRa-Prefix attribute returned in the Access-Accept message was used for IPv6 Neighbor Discovery router advertisements and the Framed-IPv6-Prefix RADIUS attribute in the Access-Accept message was used for DHCPv6 Prefix Delegation.

In this release, you can control the RADIUS IETF attribute or VSA to be used for IPv6 Neighbor Discovery router advertisements and DHCPv6 Prefix Delegation by using **aaa ipv6-nd-ra-prefix framed-ipv6-prefix** and **aaa dhcpv6-delegated-prefix delegated-ipv6-prefix** commands, respectively, in Global Configuration mode on each virtual router.

- Related Documentation**
- [Duplicate IPv6 Prefix Check Overview on page 42](#)
 - [Duplicate IPv6 Prefix Detection in the AAA User Profile Database Overview on page 42](#)

Maximum Number of IPv6 Prefixes Assigned to Clients by Using Only the DHCPv6 Local Server

You can delegate IPv6 prefixes to subscribers by using two mechanisms: ICMPv6 Neighbor Discovery Router Advertisements (NDRA) and DHCPv6 prefix delegation (PD). When the router receives the ICMPv6 router solicitation message, the DHCPv6 Solicit message, or both the messages based on the prefix advertisement mechanism, a prefix is assigned to the requesting router, which is the customer premises equipment (CPE) acting as the DHCP client at the edge of the remote client site. Consider a scenario in which the CPE device uses the PD feature alone to obtain IPv6 prefixes from the delegating router, which is the DHCPv6 local server. Also, assume that IPv6 Neighbor Discovery is not configured for allocation of prefixes to the client. In such an environment, each IPv6 subscriber uses only a single route entry and the maximum number of subscribers to which IPv6 prefixes can be delegated from the DHCPv6 local server is 64,000.

- Related Documentation**
- [Maximum Number of IPv6 Prefixes Assigned to Clients by Using Both DHCPv6 Local Server and Neighbor Discovery Router Advertisements on page 40](#)

Maximum Number of IPv6 Prefixes Assigned to Clients by Using Both DHCPv6 Local Server and Neighbor Discovery Router Advertisements

When you use both IPv6 Neighbor Discovery router advertisements and DHCPv6 Prefix Delegation methods to assign IPv6 prefixes to clients, only two host routes for IPv6 are consumed from the routing table irrespective of the way in which the router advertisement prefix is determined. The following sections describe sample configuration scenarios to illustrate how you can delegate IPv6 prefixes to a maximum of 64,000 subscribers, based on whether a unique IPv6 prefix is allocated to a client or the same IPv6 prefix is allocated to multiple clients:

- [Unique IPv6 Prefix Delegated Per Subscriber on page 41](#)

- [Same IPv6 Prefix Delegated to Multiple Subscribers on page 41](#)

Unique IPv6 Prefix Delegated Per Subscriber

Consider a scenario in which the RADIUS server is configured to assign a unique router advertisement prefix route to each IPv6 subscriber. In such a case, one route is used for neighbor discovery and one IPv6 route is consumed for prefix delegation. Therefore, it is possible to support 64,000 IPv6 bindings with 128,000 IPv6 static routes.

The following output of the **show ipv6 route** command displays how two routes are used by the same subscriber, as can be seen from the Interface field in the output. The routes are assigned using Prefix Delegation and the access-internal route, such as the DHCP and AAA/PPP host route, which is a host route to directly connected clients. Access routes, also known as AAA framed routes, are sourced by AAA.

```
host1#show ipv6 route
```

Prefix/Length	Type	Dst/Met	Interface
1111:1111:1111:1111::/64	Access	3/0	GigabitEthernet0/2.600.6
1111:1111:2222:2222:21b:c0ff:fe4	AccIntern	2/0	GigabitEthernet0/2.600.6 b:9d00/128

Same IPv6 Prefix Delegated to Multiple Subscribers

Consider a scenario in which the same prefix with a length of /64 for ICMPv6 neighbor discovery router advertisements is assigned to all subscribers by configuring the prefix in the profile or by configuring the RADIUS server to send the same prefix in the Framed-IPv6-Prefix attribute (RADIUS IETF attribute 97) of the RADIUS-Access-Accept message.

If you use this method for allocating IPv6 prefixes of /64 length to subscribers, neighbor discovery consumes one IPv6 route and prefix delegation consumes one IPv6 route, which results in a total of two IPv6 routes per subscriber being used. Therefore, it is possible to scale up to a maximum of 64,000 subscribers for delegation of IPv6 prefixes.

The increased scaling limit of support for delegation of IPv6 prefixes using the DHCPv6 local server Prefix Delegation mechanism for 64,000 subscribers applies only to E120 and E320 routers and not to ERX14xx models, ERX7xx models, and the ERX310 router because the binding information is stored in the SRP modules of E120 and E320 routers. Also, a limitation exists on the number of IPv6 interfaces and the IPv6 routing table size supported by ERX routers that prevents the support for 64,000 subscribers for Prefix Delegation on DHCPv6 local servers running on those routers.

To enable support for 64,000 subscribers for IPv6 Prefix Delegation, about 7.42 MB of memory on the SRP module is consumed additionally.

Related Documentation

- [Maximum Number of IPv6 Prefixes Assigned to Clients by Using Only the DHCPv6 Local Server on page 40](#)

Duplicate IPv6 Prefix Check Overview

You can configure AAA service to detect duplicates of IPv6 Neighbor Discovery router advertisement prefixes and DHCPv6 delegated prefixes. If a non-unique IPv6 prefix is detected by AAA, the subscriber session corresponding to the duplicate prefix is terminated.

In some network environments where the same customer logs in from multiple locations, terminating sessions with duplicate IPv6 prefixes might result in breaking subscriber setup. The duplicate IPv6 prefix-check capability is disabled by default.

If a duplicate prefix is detected by AAA before a subscriber is granted access, the subscriber is denied access. However in some cases, when two subscribers having the same IPv6 prefix log in simultaneously, the duplicate might be detected only after access is granted to both subscribers. AAA terminates the duplicate subscriber session immediately upon detecting the duplicate IPv6 prefix.



NOTE: AAA cannot detect duplicates of overlapping IPv6 prefixes.

Related Documentation

- [Configuring Duplicate IPv6 Prefix Check on page 96](#)
- [Standard RADIUS IPv6 Attributes for IPv6 Neighbor Discovery Router Advertisements and DHCPv6 Prefix Delegation Configuration on page 39](#)

Duplicate IPv6 Prefix Detection in the AAA User Profile Database Overview

You can configure AAA service to detect duplicates of both IP and IPv6 Neighbor Discovery router advertisement prefixes, Framed-IPv6-Prefixes, and DHCPv6 delegated prefixes by validating the prefixes against the AAA database instead of the IP route table. If AAA detects a non-unique IP address or IPv6 prefix, the corresponding subscriber session is terminated.

In some network environments where the same customer logs in from multiple locations, terminating sessions with duplicate IP addresses and IPv6 prefixes might result in breaking subscriber setup. The enhanced duplicate prefix detection capability is disabled by default. Because the prefix is validated against the AAA table, enabling the enhanced prefix detection capability may impact performance.

AAA maintains a new table for IPv6 prefixes and Framed-IP-Address information for subscribers. The AAA service checks for duplication of IP addresses and prefixes in this new table after PPP authorization. If a duplicate address or prefix is detected by AAA before a subscriber is granted access, the subscriber is denied access. However, in some cases, when two subscribers with the same IPv6 prefix log in simultaneously, the duplicate might be detected only after access is granted to both subscribers. AAA terminates the duplicate subscriber session immediately upon detecting the duplicate IPv6 prefix.

The following scenarios can occur during the establishment of subscriber sessions:

- When the RADIUS server assigns the same IPv6-NdRa-Prefix or Delegated-IPv6-Prefix to two subscribers, the second subscriber that contains the same prefix as the first subscriber is disconnected.
- When the RADIUS server assigns the same Framed-IPv6-Prefix to two dual-stack subscribers, the second subscriber session is rejected.
- When the RADIUS server assigns the same Framed-IP-Address and different IPv6 prefixes to two subscribers, the second subscriber session is terminated.



NOTE: AAA cannot detect duplicates of overlapping IPv6 prefixes. Also, the **aaa duplicate-prefix-check-extension** command detects duplicate prefixes globally for all VRs and is not limited to detecting duplicates on a per-VR basis.

Related Documentation

- [Configuring Detection of Duplicate IPv6 Prefixes in the AAA User Profile Database on page 97](#)
- [Monitoring Duplicate IPv6 Prefixes in the AAA User Profile Database on page 155](#)
- [Standard RADIUS IPv6 Attributes for IPv6 Neighbor Discovery Router Advertisements and DHCPv6 Prefix Delegation Configuration on page 39](#)
- *aaa duplicate-prefix-check-extension*
- *show aaa duplicate-prefix-check-extension*

Guidelines for Duplicate Address Verification

In dual-stack networks in which both IPv4 and IPv6 subscribers are available, the subscribers might be granted the same IPv4 and IPv6 addresses if one user logs in quickly after another user has logged in. To avoid the problem of two sessions containing the same address, when you enable detection of duplicate addresses, the subscriber is completely terminated when a duplicate IPv4 or IPv6 address is detected. The duplicate check operation is performed for 32-bit IPv4 subnet masks and IPv6 addresses with a prefix length of 128.

The value of the Framed-IPv6-Address attribute is determined using the Framed-IPv6-Prefix and Framed-Interface-Id attributes, normally obtained from the MAC addresses of clients in the PPP Network Control Protocol (NCP) phase in the PPP link connection process. Because the Framed-IPv6-Address attribute is not available to AAA during the authentication phase (before NCP negotiation occurs), the duplicate address detection mechanism performed for IPv4 cannot be adopted for IPv6. To achieve this functionality, if IPv6 detects a duplicate address while adding the route, it notifies AAA about the duplicate and AAA terminates the subscriber.

To correctly enable duplicate address detection when subscribers log in simultaneously, the IP and AAA applications examine the access-route table instead of the route table. In certain scenarios, AAA cannot detect whether a subscriber requesting access uses the same address as another subscriber. When the IP application detects a duplicate address

while adding the route, the IP application notifies AAA about the duplication to terminate the connection for that subscriber.

In certain cases, when two subscribers with the same address attempt to log in, the duplicate might be detected only after access is granted to both subscribers. AAA terminates the duplicate subscriber session immediately upon detecting the duplicate address.

If AAA cannot determine the virtual router (VR) context configured in the profile during subscriber authentication, the subscriber that uses the same address as another subscriber is terminated immediately after the IP application detects the duplicate address. Such a disconnection of subscribers occurs even if the duplicate subscriber was granted access previously when the VR context was not available to AAA for processing.

In a dual-stack environment in which both IPv4 and IPv6 subscribers are present, if a subscriber that uses a duplicate IPv6 address is detected, the subscriber is denied access even if the IPv4 interface address is unique. This method of terminating subscriber sessions occurs to avoid duplicate sessions from being established in scenarios in which the IPv6 interface address is the same as another client, whereas the IPv4 interface address is unique.

The following scenarios can occur during the establishment of subscriber sessions in a dual-stack network in which clients using both IPv4 and IPv6 protocols are present, and when detection of duplicate addresses is enabled on the router that delegates addresses to requesting clients. These scenarios assume that the RADIUS server is configured on a VR other than the default VR and that the AAA domain name is mapped to a non-default VR.

- When the VR context for subscribers is configured in the AAA domain map or obtained from the RADIUS server, and the same IP address is returned for two dual-stack subscribers from the RADIUS server, only the first subscriber session is configured and the second client session is terminated.
- When the same IP address is returned from the RADIUS server or the domain map for two dual-stack subscribers that log in simultaneously, only the first subscriber session is established and the second subscriber that contains the same address or prefix as the first subscriber is disconnected. Termination of the second subscriber occurs even if detection of the duplicate address occurs only after access is granted.
- When the VR context for subscribers is configured in the AAA profile, and the same IP address is returned from the RADIUS server or the domain map for two dual-stack subscribers, only the first subscriber session is configured and the second client session is terminated.
- If you disable the routing table address lookup for duplicate addresses by using the **no aaa duplicate-address-check** command, define the VR context for subscribers in the profile, and the same address is returned for two dual-stack subscribers, both the subscriber sessions are brought up successfully. However, for the second subscriber, which contains the same address as the first client, only the IPv6 interface is enabled and the IPv4 interface is not brought up.

- If the same IPv6-NdRa-Prefix (VSA 26-129) and Framed-Interface-Id (VSA 26-96) attributes are returned in the Access-Accept message from the RADIUS server for two dual-stack subscribers, and the VR context for the subscribers is specified in the profile, only the first subscriber is brought up and the second subscriber session is rejected.
- If you set the Framed-IPv6-Prefix RADIUS attribute for IPv6 Neighbor Discovery router advertisements by using the **aaa ipv6-nd-ra-prefix framed-ipv6-prefix** command, the same Framed-IPv6-Prefix (VSA 26-129) and Framed-Interface-Id (VSA 26-96) attributes are returned in the Access-Accept message from the RADIUS server for two dual-stack subscribers, and the VR context for the subscribers is specified in the profile or the domain map, only the first subscriber is brought up and the second subscriber session is rejected.
- If you set the Framed-IPv6-Prefix RADIUS attribute for IPv6 Neighbor Discovery router advertisements by using the **aaa ipv6-nd-ra-prefix framed-ipv6-prefix** command, disable the routing table address lookup for duplicate addresses, specify the VR context for subscribers in the domain map, and the same Framed-IPv6-Prefix (VSA 26-129) and Framed-Interface-Id (VSA 26-96) attributes are returned in the Access-Accept message from the RADIUS server for two dual-stack subscribers, only the first subscriber is brought up and the second subscriber session is rejected.

**Related
Documentation**

- [Configuring Duplicate IPv6 Prefix Check on page 96](#)

Propagation of LAG Subscriber Information to AAA and RADIUS

The RADIUS application sends the link aggregation group (LAG) interface ID to the RADIUS server when the subscriber is connected over LAG in DHCP standalone authenticate mode. In DHCP standalone authenticate mode, the DHCP local server enables you to configure AAA-based authentication of standalone mode DHCP clients. In addition to providing increased security, AAA authentication also provides RADIUS-based input to IP address pool selection for standalone mode clients. The RADIUS applications use the LAG interface ID to create the Acct-Session-Id, Nas-Port-Type, Nas-Port-Id, Nas-Port, and Calling-Station-Id attributes and send them to the RADIUS server in the Access-Request, Acct-Start, and Acct-Stop messages.

The RADIUS client uses one of the following LAG interface ID formats:

```
lag lag-name [.subinterface [:vlan]]
```

or

```
lag lag-name [.subinterface [:svlan-vlan]]
```

where:

- *lag-name*—Name of the LAG bundle
- *subinterface*—Number of the LAG subinterface, in the range 1–2147483647
- *vlan*—VLAN ID number
- *svlan-vlan*—S-VLAN ID number in the range 0–4095

The RADIUS application sends the LAG interface ID to the RADIUS server for all types of subscribers, such as PPP or DHCP subscribers. In this case, the LAG interface ID is displayed in the output of the **show subscribers interface** command.

The RADIUS client application creates the following RADIUS attributes based on the LAG interface ID:

[44] Acct-Session-Id—When you issue the **radius acct-session-id-format description** command, the RADIUS client uses the generic format: *erx <interface type> <interface identifier>: <hex number>* with the LAG interface ID as the interface identifier.

[61] Nas-Port-Type—When you issue the **radius ethernet-port-type** command from Global Configuration mode or the **nas-port-type ethernet** command from AAA Profile Configuration mode, RADIUS calculates the value of the Nas-Port-Type attribute. If you use neither of these commands, RADIUS uses the default [15] Nas-Port-Ethernet value for this attribute.

[5] Nas-Port—RADIUS derives a unique value from the subscriber's profileHandle and uses the value for the Nas-Port attribute. The **radius nas-port-format**, **radius vlan nas-port-format stacked**, and **radius pppoe nas-port-format** commands do not affect the value of the Nas-Port attribute.

[87] Nas-Port-Id—The **radius override nas-port-id remote-circuit-id** command configures RADIUS to use the PPPoE remote circuit ID for the Nas-Port-Id attribute. By default, RADIUS uses the LAG interface ID for the Nas-Port-Id attribute. Use the **aaa intf-desc-format include sub-intf disable** command to exclude the subinterface and S-VLAN ID in the LAG interface ID. By default, the subinterface and S-VLAN ID are included in the LAG interface ID.

[31] Calling-Station-Id—The **radius override calling-station-id remote-circuit-id** command enables RADIUS to use the PPPoE remote circuit ID for the Calling-Station-Id attribute. By default, RADIUS uses a delimited format for the interface description. The **radius calling-station-format** command does not affect the value of the Calling-Station-Id attribute.

For example, a subscriber with the default AAA or RADIUS configuration who is connected over a LAG interface lag1, with subinterface-1, VLAN ID 10, S-VLAN ID 1, and router named asterix uses the following values for RADIUS attributes in RADIUS authentication and accounting messages:

Table 10: RADIUS Attributes Specifying LAG Interface

Field Name	Field Description
Acct-Session-Id	erx lag lag1.1:1-10:0001048620
Nas-Port-Type	15
Nas-Port	2148532268
Nas-Port-Id	lag lag1.1:1-10

Table 10: RADIUS Attributes Specifying LAG Interface (*continued*)

Field Name	Field Description
Calling-Station-Id	#asterix#lag1#10

Related Documentation

- [Monitoring and Troubleshooting Remote Access](#) chapter
- [CLI Commands Used to Configure RADIUS IETF Attributes](#) on page 209
- [Configuring AAA Authentication for DHCP Local Server Standalone Mode](#) on page 497
- `show subscribers`

SRC Client Configuration Overview

The JunosE Software has an embedded client that interacts with the Juniper Networks Session and Resource Control (SRC) software, enabling the SRC software to manage the router's policy and QoS configuration.

The connection between the router and the SRC software uses the Common Open Policy Service (COPS) protocol and is fully compliant with the COPS usage for policy provisioning (COPS-PR) specification. The router's SRC client functions as the COPS client, or policy enforcement point (PEP). The SRC software functions as the COPS server, or policy decision point (PDP).

Rate limiters are aggregated for dual-stack subscribers (IPv4 and IPv6) managed by the SRC software, using external parent groups and hierarchical policy parameters. The external parent groups and policy parameters are pushed to lower interfaces from the SRC software through the Siemens Selection Switch or Service Selection Center client.



NOTE: You cannot override aggregation node values while attaching policies to the interface.

Related Documentation

- [Configuring the SRC Client](#) on page 97

SRC Client and COPS Terminology

Table 11 on page 47 provides common terms used in the COPS environment.

Table 11: SRC Client and COPS Terminology

Term	Description
COPS	Common Open Policy Service; query-and-response protocol used to exchange policy information between a policy server and its clients.
COPS-PR	COPS usage for policy provisioning; the PEP requests policy provisioning when the operational state of interface and DHCP addresses changes.

Table 11: SRC Client and COPS Terminology (*continued*)

Term	Description
PDP	Policy decision point; the COPS server, which makes policy decisions for itself and for clients that request decisions. The SRC software is the PDP.
PEP	Policy enforcement point; the COPS client, which enforces policy decisions. The JunosE COPS interface is a PEP.
PIB	Policy Information Base; a collection of sets of attributes that represent configuration information for a device.
SRC	Session and Resource Control (SRC) software, formerly the Service Deployment System (SDX) software; functions as a COPS PDP.

The JunosE Software COPS-PR implementation uses the outsourcing model that is described in RFC 3084. In this model, the PEP delegates responsibility to the PDP to make provisioning decisions on the PEP's behalf.



NOTE: When you upgrade from an earlier JunosE release, the software removes the instance of SSCC that was configured with XDR.

If you are going to perform a unified ISSU from a JunosE release numbered lower than Release 10.0.0 and you have an XDR configuration, unified ISSU is not supported while an XDR configuration is presented.

The provisioning is event-driven and is based on policy requests rather than on an action taken by an administrator—the provisioning is initiated when the PDP receives external requests and PEP events. Provisioning can be performed in bulk (for example, an entire QoS configuration) or in smaller segments (for example, updating a marking filter). The following list shows the interaction between the PEP and the PDP during the COPS-PR operation.

1. Initial connection
 - a. PEP starts the COPS-PR connection with the PDP.
 - b. PDP requests synchronization.
 - c. PEP sends all currently provisioned policies to PDP.
2. Change of interface state
 - a. PEP requests provisioning of an interface from the PDP.
 - b. PDP determines policies and sends provisioning data to the PEP.
 - c. PEP provisions the policies.
3. PDP requests policy provisioning
 - a. PDP determines new policies and sends provisioning data to the PEP.

- b. PEP provisions the policies.

The information exchange between the PDP and PEP consists of data that is modeled in Policy Information Bases (PIBs) and is encoded using the standard ASN.1 basic encoding rules (BERs).

JunosE Software uses the following PIBs:

Proprietary PIB

- JunosE-IP-PIB—This PIB defines the data model for manipulating IP service policies and addresses offered through DHCP in JunosE Software.

Non-proprietary PIBs

- COPS-PR-SPPI
- COPS-PR-SPPI-TC
- DIFFSERV-PIB
- FRAMEWORK-FEEDBACK-PIB
- FRAMEWORK-PIB
- FRAMEWORK-TC-PIB

The COPS-PR support in JunosE Software uses the proprietary PIB. This PIB consists of a series of tables that is supported in previous JunosE Software releases, including the proprietary accounting and address assignment mechanisms.

You can force the router to restart a COPS connection to, and resynchronize with, a PDP, without disabling the SRC client's COPS support. The SRC software and the SRC client maintain common state information in PIBs that both the SRC software and the SRC client use. Previously, you disabled the SRC client and reenabled it to start synchronization. The disabling of the SRC client's COPS support was undesirable for the applications that required resynchronization in addition to maintaining the COPS support. If the state of the SRC software is not synchronized with the router, the SRC software may be required to initiate resynchronization from the router.

The proprietary PIB provides the Policy Manager and QoS Manager functionality shown in the following lists.

- Policy Manager
 - Committed access rate
 - Packet filtering
 - Policy routing
 - QoS classification and marking
 - Rate limiting
 - Traffic class

- QoS Manager
 - Queues
 - Schedulers
 - Traffic classes

The JunosE-IP-PIB file is updated with each JunosE release. Since the PIB is implemented by both Juniper Networks SRC and JunosE devices, distribution of the PIB file to customers is not necessary. Customers can access the proprietary PIB file, on approval from Juniper Networks, through Juniper support.

**Related
Documentation**

- [SRC Client Configuration Overview on page 47](#)

Retrieval of DSL Line Rate Information from Access Nodes Overview

You can retrieve updated DSL line rate information from the Access Node Control Protocol (ANCP) and report this information to the SRC software with corresponding COPS messages. ANCP is also known as Layer 2 Control (L2C). To enable the router that functions as the SRC client to obtain updated line rate parameters from ANCP and transmit them to the COPS server, use the **sscc update-policy-request enable** command in Global Configuration mode. You can configure this setting on a per-virtual-router basis.

In networks with digital subscriber line access multiplexers (DSLAMs), after a connection is established between a subscriber and a routing gateway, the access node or DSLAM obtains the line rate information of the subscriber using a synchronization process. The line rate parameters are transferred in the COPS interface request by using the ANCP topology discovery message to the router that functions as the network access server (NAS). Typically, a COPS interface request is sent from the access node to the SRC client whenever an interface becomes operational.

You can configure the SRC client to obtain the line rate details from the access node whenever any change in the values of the parameters occurs. The capability to receive line rate data, when it changes on the access node, is disabled by default on the SRC client.

The access node passes the DSL line rate parameters, whenever they change, to the SRC client. The SRC client appends updated parameters to the COPS messages that it sends to the COPS server or SRC server. A COPS server processes the following topology parameters that it receives from the SRC client in the updated COPS messages:

- JunosElpInterfaceMode
- JunosElpInterfaceUpstreamRate
- JunosElpInterfaceDownstreamRate
- JunosElpInterfaceMinimumDataRateUpstream
- JunosElpInterfaceMinimumDataRateDownstream
- JunosElpInterfaceAttainableDataRateUpstream

- JunosElpInterfaceAttainableDataRateDownstream
- JunosElpInterfaceMaximumDataRateUpstream
- JunosElpInterfaceMaximumDataRateDownstream
- JunosElpInterfaceMinimumLowPowerDataRateUpstream
- JunosElpInterfaceMinimumLowPowerDataRateDownstream
- JunosElpInterfaceMaximumInterleavingDelayUpstream
- JunosElpInterfaceActualInterleavingDelayUpstream
- JunosElpInterfaceMaximumInterleavingDelayDownstream
- JunosElpInterfaceActualInterleavingDelayDownstream
- JunosElpInterfaceDSLlinestate

A COPS server that runs an SRC software release earlier than Release 3.0.0 does not support and process the preceding topology parameters that are appended to the COPS messages. Such COPS servers analyze the information, other than the parameters that describe updated DSL line rate details, that they receive in the COPS messages for policy management. Therefore, the COPS-PR operation ensures backward compatibility of the SRC clients with the COPS servers running SRC software releases earlier than Release 3.0.0 by ignoring the received line rate details.

When you configure the **sscc update-policy-request enable** command, a warning message is displayed, prompting you to confirm whether you want to enable the router that functions as the SRC client to forcibly send line rate information parameters to the COPS server, which is running a release of SRC software earlier than Release 3.0.0 that is not compatible with the line rate message format.

Even if you confirm the prompt to enable the SRC client to forcibly send updated DSL line rate parameters to the COPS server, the COPS server that is running a release of SRC software earlier than Release 3.0.0 ignores the updated line rate details that it receives and processes only the other information in the COPS messages.

The Policy Information Base (PIB) is modified to extend the JunosElpInterfaceEntry object. ANCP now notifies the SRC software about any change in the ANCP parameters. If this change in rate is greater than 10 percent or a change in mode, SRC software reports this upgrade to the service activation engine (SAE) in SRC version 3.0.0 and later.

Related Documentation

- [SRC Client Configuration Overview on page 47](#)
- [Monitoring the SRC Client Connection Status on page 155](#)
- *sscc update-policy-request enable*

DHCPv6 Local Address Pools for Allocation of IPv6 Prefixes Overview

In previous releases, you configured DHCPv6 local servers on a virtual router to delegate IPv6 prefixes to DHCPv6 clients. In this release, you can configure IPv6 local address pools to allocate IPv6 prefixes to clients in networks that use DHCPv6. These pools can

be used to assign prefixes from a delegating router, which is an E Series router configured as a DHCPv6 local server, to the requesting router, which is the customer premises equipment (CPE) at the edge of the remote client site that acts as the DHCP client.

The DHCPv6 prefix delegation feature is useful in scenarios in which the delegating router does not have information about the topology of the networks in which the customer edge device or requesting router is located. In such cases, the delegating router requires only the identity of the requesting router to choose a prefix for delegation. An IPv6 local pool is configured on the delegating router, which contains information about the prefixes, their validity periods, and other parameters to control their assignment to the requesting routers. The delegating router is configured with a set of prefixes that is used to assign to a CPE or DHCPv6 client, when it first establishes a connection with an Internet service provider (ISP).

When the delegating router receives a request from a DHCPv6 client, it selects an available prefix and delegates it to the client. The DHCPv6 client subnets the delegated prefix and assigns the prefixes to links at the customer edge.

Keep the following points in mind when you configure IPv6 local address pools to assign prefixes to requesting routers:

- You must enable the IPv6 local address pool feature to be able to configure IPv6 local address pools.
- You can configure IPv6 local address pools for DHCP to allocate prefixes to client requests that are received over PPP or non-PPP links, such as VLAN, S-VLAN, or Ethernet.
- You can configure multiple local address pools on a single virtual router, up to a maximum of 500 pools per virtual router.
- You can also configure multiple address pools on multiple virtual routers. Each IPv6 local address pool must have a unique name.
- You can configure a valid and preferred lifetime for each IPv6 prefix, which determines the length of time the requesting router can use the prefix.
- You can configure multiple prefix ranges in an IPv6 local pool. The ranges can have the same or different assigned prefix lengths.
- You cannot configure overlapping prefix ranges in an IPv6 local pool. If you try to configure a prefix range that overlaps with an existing prefix range in the IPv6 local pool, an error message is displayed stating that the prefix range could not be configured. Similarly, an error message is displayed if you try to configure a prefix range in an IPv6 local pool that overlaps with a prefix range in another IPv6 local pool on the same virtual router.
- You can configure certain prefix ranges to be excluded from being used for delegation to the requesting router.
- You can configure the IPv6 addresses of a primary and secondary DNS server in an IPv6 local pool. The DNS server addresses are returned to the client in DHCPv6 responses as part of the DNS Recursive Name Server option.

- You can configure a list of up to four domain names in an IPv6 local pool to be used during the resolution of hostnames to IP addresses. These domain names are returned to clients in the DHCPv6 responses as part of the Domain Search List option.
- You can configure an IPv6 local address pool in an AAA domain map to assign prefixes to requesting DHCPv6 clients using the **ipv6 prefix-pool-name** command in Domain Map Configuration mode. If the authentication server returns the IPv6 local address pool name in the Framed-IPv6-Pool attribute of the RADIUS-Access-Accept message, this pool overrides the IPv6 local address pool configured in the domain map.
- You cannot delete a pool or a prefix range from which prefixes have been allocated to requesting routers or DHCPv6 clients. However, you can forcibly delete such a pool or prefix range by using the **force** keyword in the **ipv6 local pool poolName** and **prefix** commands. If a pool is deleted or the prefix range associated with the pool is deleted, and prefixes have been assigned to DHCPv6 clients or requesting routers, the corresponding DHCPv6 bindings are also deleted.
- When multiple prefix ranges are configured in a pool, the DHCPv6 prefix delegation feature allocates prefixes from the configured ranges in the order of the assigned prefix length. The delegating router or the DHCPv6 server attempts to allocate a prefix from the range with lowest assigned prefix length. If this attempt fails because the pool has been fully allocated, the server tries to allocate a prefix from the subsequent prefix ranges. These ranges could have the same prefix length as the first one or a higher length.



NOTE: Although you can configure an IPv6 local pool with the assigned prefix length as /128, which implies a full IPv6 address, this assignment is not useful for the DHCPv6 prefix delegation feature because it assigns a prefix with a length of only /64 or less. A pool with an assigned prefix length of /128 is useful when complete IPv6 addresses are assigned to the DHCPv6 clients.

- When an IPv6 client that is connected to the requesting router using a PPP link is delegated a prefix by the DHCPv6 server, the client binding is removed when the PPP interface goes down and is not retained until the lease time expires. A new client binding is created for the PPP subscriber in response to a renew or rebind request sent to the DHCP server. This method of re-creating the client binding ensures that the client receives a new authentication configuration and is assigned a prefix when it sends a rebind or renew request after the PPP interface flaps (constantly goes up and down).

When a PPP user establishes a PPP connection with the E Series router functioning as a remote access server, the subscriber is first authenticated using the RADIUS protocol. The Access-Accept message returned from the RADIUS server can contain different IPv6 attributes, including the Framed-IPv6-Pool attribute, which contains the name of the IPv6 pool from which a prefix needs to be assigned to the subscriber. The prefix is assigned to the subscriber using the DHCPv6 prefix delegation feature, which is covered in the next section.

Related Documentation

- [Example: Delegating the DHCPv6 Prefix on page 54](#)

Example: Delegating the DHCPv6 Prefix

Consider a scenario in which a number of devices on a home network are connected to a customer premises equipment, CPE1, which is the requesting router. CPE1 is connected using a PPP link to the provider edge device, PE1, which is an E Series router operating as the DHCPv6 server or delegating router. After the IPv6 link is formed between CPE1 and PE1 and the IPv6 link-local address is created, CPE1 requests and obtains prefixes that are shorter than /64 (usually of length, /48) from PE1.

CPE1 is connected to the home network. CPE1 divides the single delegated prefix that it received from PE1 into multiple /64 prefixes and assigns one /64 prefix to each of the links in the home network. The address allocation mechanism in the subscriber network can be performed using ICMPv6 Neighbor Discovery in router advertisements, DHCPv6, or a combination of these two methods.

When PE1 receives a request for prefix delegation from CPE1, PE1 assigns prefixes from the list of unallocated prefixes in the IPv6 local pool.

The following sections of this example show how to delegate the DHCPv6 prefix:

- [Order of Preference in Determining the Local Address Pool for Allocating Prefixes on page 54](#)
- [Order of Preference in Allocating Prefixes and Assigning DNS Addresses to Requesting Routers on page 55](#)

Order of Preference in Determining the Local Address Pool for Allocating Prefixes

You can configure multiple local address pools on a virtual router. When multiple pools are configured, the pool that is used to allocate the prefix to the requesting router is selected using the following order of preference:

- If a pool name is returned by the RADIUS server in the Framed-IPv6-Pool attribute or in the Delegated-Ipv6-Pool attribute (VSA 26-161), that pool is used to delegate the prefix to the client.
- If the **aaa dhcpv6-ndra-pool override** command is not configured, and if the RADIUS server returns a pool name in the Framed-IPv6-Pool attribute, that pool name is used to delegate the prefix to the client.
- If the **aaa dhcpv6-ndra-pool override** command is configured, and if the RADIUS server returns a pool name in the Delegated-Ipv6-Pool attribute (VSA 26-161), that pool name is used to delegate the prefix to the client.
- If the RADIUS server does not return the pool name, the pool name configured in the AAA domain map (Ipv6-Prefix-Pool-Name) is used to delegate the prefix to the client.
- If no local address pool name is configured in the AAA domain map, the IPv6 address of the interface on which the request was received is used to determine the pool.
- If the interface address matches with any of the prefix ranges configured in the IPv6 local address pool on the router, that pool is used to delegate the prefix to the client.

Order of Preference in Allocating Prefixes and Assigning DNS Addresses to Requesting Routers

Prefix delegation can be configured at the interface level and at the router level. Also, certain VSA attributes returned in the RADIUS Access-Accept message from the authentication server can impact the selection of the prefix to be assigned to the requesting router. The level of preference attached to each of these prefix delegation configurations is crucial. The delegating router uses the following order of preference to determine the source from which the DHCPv6 prefix is delegated to the requesting router from the DHCPv6 server:

1. An interface that is configured for prefix delegation is given priority over the RADIUS attributes returned in the Access-Accept message or the prefixes configured in the IPv6 local address pool on the delegating router.
2. The RADIUS server might return one or more of the following attributes in the Access-Accept message in response to the client authentication request:
 - Ipv6-NdRa-Prefix (VSA 26-129)
 - Framed-IPv6-Prefix (RADIUS IETF attribute 97)
 - Delegated-IPv6-Prefix (RADIUS IETF attribute 123)
 - Framed-IPv6-Pool (RADIUS IETF attribute 100)
 - Delegated-Ipv6-Pool (VSA 26-161)

If any of the first three attributes are returned, then the prefix contained in those attributes is used and the pool name in the Framed-IPv6-Pool/Delegated-Ipv6-Pool attribute is ignored. For example, if both the Delegated-IPv6-Prefix or Framed-IPv6-Prefix, and Framed-IPv6-Pool/Delegated-Ipv6-Pool attributes are returned from the RADIUS server, the DHCPv6 prefix delegation mechanism uses the Delegated-IPv6-Prefix attribute to advertise the prefix to clients.

3. If prefix delegation is not configured at the interface level and if no prefix is returned from the attribute in the RADIUS Access-Accept message, the prefix configured in the IPv6 local pool is delegated to the requesting router.

If you configured a list of IPv6 DNS servers and a string of domain names in the IPv6 local address pool, the order of preference in returning the DNS server address or domain name to the requesting client in the DHCPv6 response is as follows:

- Information returned from the RADIUS server for DNS servers only
- Information from the pool
- Locally configured DNS attributes

Related Documentation

- [Example: Limiting the Number of Prefixes Used by DHCPv6 Clients on page 103](#)
- [DHCPv6 Local Address Pools for Allocation of IPv6 Prefixes Overview on page 51](#)

Sending of Delegated-IPv6-Prefix from the DHCPv6 Local Server to SRC Client Overview

You can configure a DHCPv6 local server to send the Delegated-IPv6-Prefix with a PPPv6 interface name to an SRC client after delegating IPv6 prefixes to a CPE. The SRC client sends the Delegated-IPv6-Prefix to an SRC server. The SRC server uses the Delegated-IPv6-prefix to authenticate DHCPv6 subscribers logging in to the website for service activation. The SRC server authenticates the DHCPv6 subscribers by comparing their IPv6 prefixes with the set of IPv6 prefixes stored in the SRC server. After authentication, the DHCPv6 subscriber activates a service. On service activation, the SRC server pushes the respective service classifiers to the PPPv6 interface.



NOTE:

- Because an interface supports only 512 classifiers per direction (ingress or egress), there is a possibility to have subscribers without any service.
- For all subscribers, the SRC server sends the Acct-Stop message to the RADIUS server only on CPE termination or DHCPv6 prefix release, but not when the subscriber is terminated.
- The DHCPv6 local server sends the Delegated-IPv6-prefix allocation or removal notifications to an SRC client only when the DHCP subscribers are in bound state.

You can enable the DHCPv6 local server to send the Delegated-IPv6-Prefix with the PPPv6 interface name to the SRC client by using the ***ipv6 dhcpv6-local ssc*** command. Before you enable this feature, you must enable the SRC client's COPS support by using the ***sscc enable*** command, the policy and QoS configuration support for IPv6 interfaces by using the ***sscc protocol ipv6*** command, and the DHCPv6 local server by using the ***service dhcpv6-local*** command.

This feature does not support the following:

- Dynamic subscriber interface for DHCPv6 subscribers or DHCPv6 prefix delegation subscribers.
- SRC-based authentication for DHCPv6 prefix delegation over PPPv6.
- Policy-based service activation due to unavailability of dynamic subscriber interface.

The COPS-PR Policy Information Base (PIB) is modified with a new class ***junoseipv6DhcpAddressClasses*** and two new objects such as ***Junoseipv6AddressRequestEntry*** and ***Junoseipv6DhcpDecisionEntry*** for the following:

- To support communication synchronization between JunosE router and SRC server.
- To exchange the delegated IPv6 prefix and its information between BRAS and SRC server.
- To exchange the information between DHCPv6 plug-in in SRC client and SRC server by using REQ and DEC messages.

- Related Documentation**
- [Configuring the DHCPv6 Local Server on page 501](#)
 - [Monitoring the DHCPv6 Local Server to Send Delegated-IPv6-Prefix to SRC Client on page 602](#)
 - `show ipv6 dhcpv6-local ssc`

IPv6 Prefix Allocation Using Neighbor Discovery Router Advertisements from IPv6 Address Pools Overview

You can configure IPv6 local address pools for Neighbor Discovery router advertisements on a virtual router in order to allocate prefixes to Neighbor Discovery clients. These pools can be used to assign prefixes from the E Series router.

An IPv6 local address pool for Neighbor Discovery router advertisements is configured on the router running the B-RAS application, which contains information about the prefixes. When the B-RAS application running on the E Series router receives a request from a PPP IPv6 client, it selects an available prefix and allocates it to the client.

Allocation of Neighbor Discovery Prefixes for IPv6 Subscribers over PPP Links

When a PPP user establishes a PPP connection with the E Series router functioning as a remote access server, the subscriber is first authenticated using the RADIUS protocol. The Access-Accept message returned from the RADIUS server can contain different IPv6 attributes, including the IPv6-NdRa-Pool attribute, which contains the name of the IPv6 pool from which a prefix needs to be assigned to the subscriber. The prefix is assigned to the subscriber using the Neighbor Discovery router advertisements feature.

Order of Preference in Determining the Local Address Pool for Allocating Prefixes for Neighbor Discovery Router Advertisements

You can configure multiple local address pools for Neighbor Discovery router advertisements on a virtual router. When multiple pools are configured, the pool that is used to allocate the prefix to the requesting PPPv6 subscriber is selected using the following order of preference:

1. If the **aaa dhcpv6-ndra-pool override** command is not configured and a pool name is returned by the RADIUS server in the IPv6-Ndra-Pool attribute, that pool is used to allocate the prefix to the client.
2. If the **aaa dhcpv6-ndra-pool override** command is configured and a pool name is returned by the RADIUS server in the Framed-Ipv6-Pool attribute, that pool is used to allocate the prefix to the client.
3. If the RADIUS server does not return a pool name in either of the above-mentioned points, based on the **aaa dhcpv6-ndra-pool override** command, the pool name configured in the AAA domain map is used.

Order of Preference in Assigning Prefixes when Neighbor Discovery Router Advertisements are Configured on an Interface

The router running the B-RAS application uses the following order of preference to determine the source from which the Neighbor Discovery router advertisements prefix is allocated to the requesting PPPv6 subscriber from the Neighbor Discovery Router Advertisements server:

1. An interface that is configured for the Neighbor Discovery router advertisements prefix is given priority over the RADIUS attributes returned in the Access-Accept message or the prefixes configured in the IPv6 local address pool for Neighbor Discovery router advertisements on the router running the B-RAS application.
2. The RADIUS server might return one or more of the following attributes in the Access-Accept message in response to the client authentication request:
 - Ipv6-NdRa-Prefix (VSA 26-129)
 - Framed-IPv6-Prefix (RADIUS IETF attribute 97)
 - Framed-IPv6-Pool (RADIUS IETF attribute 100)
 - IPv6-Ndra-Pool (VSA 26-157)

If either of the first two attributes are returned, then the prefix contained in those attributes is used, and the pool name in the Framed-IPv6-Pool or Ipv6-Ndra-Pool attribute is ignored.

3. If the RADIUS server does not return any of the above-mentioned attributes, the IPv6 prefix pool name of the Neighbor Discovery router advertisements mentioned in the AAA domain map will be used to allocate the prefix to the requesting PPPv6 subscriber.

Guidelines for Allocating Neighbor Discovery Prefixes Using IPv6 Address Pools

The following are guidelines for allocating prefixes using IPv6 address pools for Neighbor Discovery router advertisements:

- You must enable the IPv6 local address pool for the Neighbor Discovery router advertisements feature to be able to configure IPv6 local address pools for Neighbor Discovery router advertisements.
- You can configure IPv6 local address pools for Neighbor Discovery router advertisements to allocate prefixes to client requests that are received over PPP.
- You can configure multiple local address pools on a single virtual router up to a maximum of 500 pools per virtual router.
- You can also configure multiple address pools on multiple virtual routers. Each IPv6 local address pool must have a unique name.
- You can configure up to ten prefix ranges in an IPv6 local address pool. The ranges can have only /64 prefix length.
- You can configure a maximum of 1,048,576 prefixes per prefix range to be used for allocation of prefixes to clients using Neighbor Discovery router advertisements. If you

attempt to configure prefixes after the maximum limit of prefixes per prefix range is exceeded, a warning message stating that automatic truncation will be performed is displayed.

- You can configure a maximum of 400,000,000 prefixes throughout the system for allocation of prefixes using Neighbor Discovery router advertisements. An error message is displayed if you attempt to configure a prefix for a pool when this maximum system-wide limit is exceeded.
- If you configure the maximum number of IPv6 prefixes, which is 1,048,576 per prefix range, for the first 383 local address pools for Neighbor Discovery router advertisements by using the **ipv6 local ndra-pool *poolName*** command, the system-wide maximum limitation of 400,000,000 is reached. In such a case, if you attempt to configure the IPv6 prefix ranges to be allocated for the 384th pool, an error message is displayed stating that the prefix cannot be configured. Although all of the 500 IPv6 local address pools are configured correctly, you cannot configure prefixes for Neighbor Discovery from the 384th pool through the 500th pool because the maximum number of prefixes supported for the entire system is reached with the 383rd pool.
- You cannot configure overlapping prefix ranges in an IPv6 local pool. If you try to configure a prefix range that overlaps with an existing prefix range in the IPv6 local pool, an error message is displayed stating that the prefix range could not be configured. Similarly, an error message is displayed if you try to configure a prefix range in an IPv6 local pool that overlaps with a prefix range in another IPv6 local pool on the same virtual router.
- You can configure certain prefix ranges to be excluded from being used for allocation to the requesting subscriber.
- You can configure the name of an IPv6 local address pool in an AAA domain map using the **ipv6-ndra-pool-name** command in Domain Map Configuration mode. If the authentication server returns the IPv6 local address pool name in the Framed-IPv6-Pool attribute or Ipv6-NdRa-Pool attribute of the RADIUS-Access-Accept message, this pool overrides the IPv6 local address pool configured in the domain map.
- You cannot delete a pool or a prefix range from which prefixes have been allocated to requesting routers or Neighbor Discovery router advertisements clients. However, you can forcibly delete such a pool or prefix range by using the **force** keyword in the **ipv6 local ndra-pool *poolName*** and **ndraprefix** commands. If a pool is deleted or the prefix range associated with the pool is deleted forcibly, corresponding subscribers will be logged out forcibly.
- Two new RADIUS attributes are added: Ipv6-Ndra-Pool and Delegated-Ipv6-Pool. For more information on these attributes see [“Juniper Networks VSAs” on page 248](#).
- You can issue the **aaa dhcpv6-ndra-pool override** command to use Framed-Ipv6-Pool attribute for IPv6 Neighbor Discovery router advertisements and the Delegated-Ipv6-Pool attribute for DHCPv6 Prefix Delegation. The **no** version of this command causes the Ipv6-NdRa-Pool attribute to be used for IPv6 Neighbor Discovery router advertisements and the Framed-Ipv6-Pool attribute to be used for DHCPv6 Prefix Delegation.
- If you want the IPv6-NdRa-Prefix attribute to be included in the Acct-Start messages that the router sends to the RADIUS server, you can use the **radius include**

ipv6-ndra-prefix acct-start enable command. In such a case, the prefix allocated to the subscriber from the IPv6 local address pool for Neighbor Discovery is included in the Ipv6-NdRa-Prefix attribute or the Framed-Ipv6-Prefix attribute.

Similarly, to cause the Ipv6-NdRa-Prefix attribute to be included in the Acct-Stop messages sent to the RADIUS server, you can use the **radius include ipv6-ndra-prefix acct-stop enable** command. You can use the **disable** keyword with the **radius include ipv6-ndra-prefix acct-start** and **radius include ipv6-ndra-prefix acct-stop** commands to prevent the Ipv6-NdRa-Prefix attribute to be sent in the Acct-Start or Acct-Stop messages.

**Related
Documentation**

- [Configuring the DHCPv6 Local Address Pools on page 100](#)
- [Configuring IPv6 Neighbor Discovery Local Address Pools on page 105](#)
- *aaa dhcpv6-ndra-pool override*
- *ipv6 address-pool ndra*
- *ipv6 local ndra-pool*

Understanding IPCP and IPv6CP Negotiations for IPv4 and IPv6 Clients Based on RADIUS-Returned Attributes

Point-to-Point Protocol (PPP) uses Internet Protocol Control Protocol (IPCP) and Internet Protocol version 6 Control Protocol (IPv6CP) negotiations for assigning IP version 4 (IPv4) and IP version 6 (IPv6) addresses to authenticated PPP Broadband Remote Access Server (B-RAS) subscribers by using RADIUS-returned attributes or the local address pool configured on the router.

You can now enable IPCP and IPv6CP negotiations for IPv4 and IPv6 clients based only on the RADIUS attributes present in the Access-Accept message returned from the RADIUS server.

IPCP negotiation is initiated for IPv4 clients only when the Framed-Ip-Address [8] attribute or Framed-Pool [88] attribute is returned from the RADIUS server. IPv6CP negotiation is initiated for IPv6 clients only when the Framed-Interface-Id [96] attribute, IPv6 prefix attributes, or IPv6 pool name attributes are returned from the RADIUS server.

You can issue the **aaa radius-override-ncp-negotiation** command with the **enable** keyword to enable IPCP and IPv6CP negotiations based on RADIUS-returned attributes.



NOTE: You can enable IPCP and IPv6CP negotiations based on RADIUS-returned attributes only for Point-to-Point Protocol (PPP) Broadband Remote Access Server (B-RAS) subscribers.

When you enable IPCP and IPv6CP negotiations based on RADIUS-returned attributes, IPCP negotiation is initiated for IPv4 clients only if one of the following conditions is satisfied:

- Framed-Ip-Address [8] attribute is returned from the RADIUS server in the Access-Accept message.
- Framed-Pool [88] attribute is returned from the RADIUS server in the Access-Accept message and the B-RAS successfully allocates an IP address to a PPP B-RAS subscriber from the received pool.

When you enable IPCP and IPv6CP negotiations based on RADIUS-returned attributes, IPv6CP negotiation is initiated for IPv6 clients only if one of the following conditions is satisfied:

- Framed-Ipv6-Prefix [97] attribute is returned from the RADIUS server in the Access-Accept message and the **aaa ipv6-nd-ra-prefix framed-ipv6-prefix** command is configured.
- Ipv6-NdRa-Prefix [26-129] VSA is returned from the RADIUS server in the Access-Accept message and the **aaa ipv6-nd-ra-prefix framed-ipv6-prefix** command is not configured or the **no aaa ipv6-nd-ra-prefix framed-ipv6-prefix** command is configured.
- Framed-Ipv6-Pool [100] attribute is returned from the RADIUS server in the Access-Accept message, the **aaa dhcpv6-ndra-pool override** command is configured, and the B-RAS successfully allocates a prefix to a PPP B-RAS subscriber from the IPv6 neighbor discovery route advertisement pool for neighbor discovery.
- Ipv6-Ndra-Pool [26-157] VSA is returned from the RADIUS server in the Access-Accept message, the **aaa dhcpv6-ndra-pool override** command is not configured or the **no aaa dhcpv6-ndra-pool override** command is configured, and the B-RAS successfully allocates a prefix to a PPP B-RAS subscriber from the IPv6 neighbor discovery route advertisement pool for neighbor discovery.
- Framed-Interface-Id [96] attribute is returned from the RADIUS server in the Access-Accept message.
- Delegated-Ipv6-Prefix [123] attribute is returned from the RADIUS server in the Access-Accept message.

You can issue the **aaa radius-override-ncp-negotiation** command with the **disable** keyword or the **no aaa radius-override-ncp-negotiation** command to enable IPCP and IPv6CP negotiations based on RADIUS-returned attributes and configurations done on the subscriber domain map, which is the default behavior of the router.

When IPCP and IPv6CP negotiations are based on RADIUS-returned attributes and configurations done on the subscriber domain map, IPCP negotiation is initiated for IPv4 clients only if one of the following conditions is satisfied:

- Framed-Ip-Address [8] attribute is returned from the RADIUS server in the Access-Accept message.
- Framed-Pool [88] attribute is returned from the RADIUS server in the Access-Accept message and the B-RAS successfully allocates an IP address to a PPP B-RAS subscriber from the received pool.

- You have issued the **address-pool-name** command to configure an address pool name on the subscriber domain map and the B-RAS successfully allocates an IP address to a PPP B-RAS subscriber from the configured pool name.
- You have configured an IP local pool for the subscriber virtual router and the B-RAS allocates an IP address to a PPP B-RAS subscriber from the configured pool.

When IPCP and IPv6CP negotiations are based on RADIUS-returned attributes and configurations done on the subscriber domain map, IPv6CP negotiation is initiated for IPv6 clients only if one of the following conditions is satisfied:

- Framed-Ipv6-Prefix [97] attribute is returned from the RADIUS server in the Access-Accept message.
- Ipv6-NdRa-Prefix [26-129] VSA is returned from the RADIUS server in the Access-Accept message.
- Framed-Ipv6-Pool [100] attribute is returned from the RADIUS server in the Access-Accept message.
- Ipv6-Ndra-Pool [26-157] VSA is returned from the RADIUS server in the Access-Accept message.
- Delegated-Ipv6-Pool [26-161] VSA is returned from the RADIUS server in the Access-Accept message.
- Delegated-Ipv6-Prefix [123] attribute is returned from the RADIUS server in the Access-Accept message.
- You have issued the **ipv6-prefix-pool-name** command to configure an IPv6 local prefix pool on the subscriber domain map.
- You have issued the **ipv6-ndra-pool-name** command to configure an IPv6 neighbor discovery router advertisement pool on the subscriber domain map and the B-RAS allocates a neighbor discovery router advertisement prefix to a PPP B-RAS subscriber from the configured pool.

Related Documentation

- *Understanding PPP*
- [Overriding AAA to Perform IPCP and IPv6CP Negotiations Based on RADIUS-Returned Attributes on page 107](#)
- *aaa dhcpv6-ndra-pool override*
- *aaa ipv6-nd-ra-prefix framed-ipv6-prefix*
- *aaa radius-override-ncp-negotiation*
- *address-pool-name*
- *ipv6-ndra-pool-name*
- *ipv6-prefix-pool-name*

Sending the Acct-Start and Acct-Update Messages After NCP Negotiation Is Completed Overview

When authentication, authorization, and accounting (AAA) receives authentication-grant permission from the RADIUS server, an Acct-Start message is sent immediately from the Broadband Remote Access Server (B-RAS) to the RADIUS server. An Acct-Update message is sent from the B-RAS to the RADIUS server immediately after IP notifies AAA to set a user IP address or user interface ID.

However, you can use the **aaa accounting delay-start** command to delay the sending of the Acct-Start message from the B-RAS to the RADIUS server until Network Control Protocol (NCP) negotiation is completed. When NCP negotiation is completed, IP notifies AAA to set a user IP address or user interface ID (in case of Point-to-Point Protocol version 6 [PPPoV6] subscribers). Only then does the B-RAS send the Acct-Start message to the RADIUS server. This ensures that Framed-Interface-Id [96] or Framed-IP-Address [8] RADIUS attributes that are negotiated during Internet Protocol Control Protocol (IPCP) or Internet Protocol version 6 Control Protocol (IPv6CP) are present in the Acct-Start message, thereby reducing the number of accounting messages from two to one for each subscriber.

When delaying the sending of the Acct-Start message is enabled:

- In a single-stack environment:
 - Acct-Start message is sent after IPCP negotiation is completed (for PPPv4 subscribers)
 - Acct-Start message is sent after IPv6CP negotiation is completed (for PPPv6 subscribers)
- In a dual-stack environment:
 - Acct-Start message is sent after IPCP or IPv6CP negotiation is completed, whichever occurs first (for PPPv4 and PPPv6 subscribers)

You can use the **aaa accounting immediate-update-framed-ipv6-interfaceId-negotiation** command to send the Acct-Update message with the Framed-Interface-Id [96] RADIUS attributes from the B-RAS to the RADIUS server for PPPv6 users immediately after NCP negotiation is completed.

When the immediate sending of the Framed-Interface-Id [96] RADIUS attributes in the Acct-Update message is enabled:

- In a single-stack environment:
 - Acct-Update message is sent after IPv6CP negotiation is completed (for PPPv6 subscribers)
- In a dual-stack environment:

- Acct-Update message is sent after both IPCP and IPv6CP negotiations are completed (for PPPv4 and PPPv6 subscribers)

**Related
Documentation**

- [Sending the Acct-Start and Acct-Update Messages After NCP Negotiation Is Completed on page 108](#)
- [Monitoring AAA Accounting Configuration on page 115](#)
- [Monitoring AAA Configuration to Send the Acct-Start Message After NCP Negotiation Is Completed on page 117](#)
- [Monitoring AAA Configuration to Send the Acct-Update Message After NCP Negotiation Is Completed on page 117](#)
- *aaa accounting delay-start*
- *aaa accounting immediate-update-framed-ipv6-interfaceId-negotiation*

CHAPTER 2

Configuring Remote Access

- Remote Access Configuration Tasks on page 66
- Configuring a B-RAS License on page 67
- Configuring AAA Duplicate Accounting on page 68
- Configuring AAA Broadcast Accounting on page 68
- Overriding AAA Accounting NAS Information on page 68
- Collecting Accounting Statistics on page 69
- Configuring RADIUS AAA Servers on page 69
- Configuring SNMP Traps on page 71
- Creating the AAA Local Authentication Environment on page 73
- Creating AAA Local User Databases on page 73
- Adding AAA User Entries to Local User Databases on page 74
- Adding AAA User Entries to Default Local User Databases on page 74
- Configuring AAA User Entries in Local User Databases on page 75
- Assigning a Local User Database to a Virtual Router on page 76
- Enabling Local Authentication on the Virtual Router on page 77
- Example: Configuring AAA Local Authentication on page 77
- Configuring DNS Primary and Secondary NMS on page 81
- Configuring WINS Primary and Secondary NMS on page 81
- Configuring a Local Address Server on page 81
- Creating an IP Interface on page 83
- Controlling Access to Domain Names on page 85
- Example: Associating all Subscribers of a PPP Interface with a Specific Domain Name on page 86
- Example: Associating Multiple Domain Names with a Specific Domain Name on page 87
- Configuring an AAA Per-Profile Attribute List on page 88
- Configuring Interim, Broadcast, and Policy-Based Accounting in Virtual Router Groups and PPP Profiles on page 89
- Configuring the NAS-Port-Type Attribute Manually on page 91

- [Configuring a Service Description for the AAA Profile on page 91](#)
- [Configuring the Route-Download Server to Download Routes on page 92](#)
- [Configuring the Router to Obtain the LLID for a Subscriber on page 93](#)
- [Troubleshooting Subscriber Preauthentication on page 95](#)
- [Configuring Custom Mappings for PPP Terminate Reasons on page 95](#)
- [Configuring Duplicate IPv6 Prefix Check on page 96](#)
- [Configuring Detection of Duplicate IPv6 Prefixes in the AAA User Profile Database on page 97](#)
- [Configuring the SRC Client on page 97](#)
- [Configuring the Forwarding of COPS Requests to the SRC Server Based on DCM Profiles on page 100](#)
- [Configuring the DHCPv6 Local Address Pools on page 100](#)
- [Example: Limiting the Number of Prefixes Used by DHCPv6 Clients on page 103](#)
- [Example: Using DHCPv6 Local Address Pools for Prefix Delegation over non-PPP Links on page 104](#)
- [Configuring IPv6 Neighbor Discovery Local Address Pools on page 105](#)
- [Overriding AAA to Perform IPCP and IPv6CP Negotiations Based on RADIUS-Returned Attributes on page 107](#)
- [Sending the Acct-Start and Acct-Update Messages After NCP Negotiation Is Completed on page 108](#)

Remote Access Configuration Tasks

Before you begin to configure B-RAS, you need to collect the following information for the RADIUS authentication and accounting servers:

- IP addresses
- User Datagram Protocol (UDP) port numbers
- Secret keys

Each configuration task is presented in a separate section in this chapter. Most of the B-RAS configuration tasks are optional.

To configure B-RAS, perform the following tasks:

1. Configure a B-RAS license.
2. (Optional) Map a user domain name to a virtual router. By default, all requests go through a default router.
3. (Optional) Set up domain name and realm name usage.
4. (Optional) Specify a single name for users from a domain.
5. Configure an authentication server on the router.
6. (Optional) Configure UDP checksums.

7. (Optional) Configure an accounting server on the router.
8. (Optional) Configure Domain Name System (DNS) and Windows Internet Name Service (WINS) name server addresses.
9. (Optional) Configure a local address pool for remote clients.
10. (Optional) Configure one or more DHCP servers.
11. Create a PPP interface on which the router can dynamically create an IP interface.
12. (Optional) Configure AAA profiles.
13. (Optional) Use vendor-specific attributes (VSAs) for Dynamic Interfaces.
14. (Optional) Set idle or session timeout.
15. (Optional) Limit the number of active subscribers on a virtual router (VR) or port.
16. (Optional) Set up the router to notify RADIUS when a user fails AAA.
17. (Optional) Configure a RADIUS download server on the router.
18. (Optional) Configure the Session and Resource Control (SRC) client (formerly the SDX client).
19. (Optional) Set baselines for AAA statistics or RADIUS authentication and accounting statistics.

Related Documentation • [Remote Access Overview on page 4](#)

Configuring a B-RAS License

From Global Configuration mode, configure a B-RAS license:

```
host1(config)#license b-ras k3n91s6gvtj
```

B-RAS licenses are available in various sizes to enable subscriber access for up to one of the following maximum number of simultaneous active IP, LAC, and bridged Ethernet interfaces:

- 4000
- 8000
- 16,000
- 32,000
- 48,000



NOTE: To use a B-RAS license for 16,000 or more interfaces, each of your SRP modules must have 1 gigabyte (GB) of memory.

Related Documentation

- *license b-ras*

Configuring AAA Duplicate Accounting

To configure and enable duplicate accounting on a virtual router, you use the **aaa accounting duplication** command with the virtual router name to which AAA information is sent to the accounting server on that virtual router. For example, to enable duplicate accounting for the default virtual router:

```
host1(config)#aaa accounting duplication xyzCompanyServer
```

Related Documentation

- *aaa accounting duplication*

Configuring AAA Broadcast Accounting

To configure and enable broadcast accounting on a virtual router:

1. Create the virtual router group and enter VR Group Configuration mode:

```
host1(config)#aaa accounting vr-group groupXyzCompany
host1(vr-group-config)#
```

2. Add up to four virtual routers to the group. The accounting information will be sent to all virtual routers in the group.

```
host1(vr-group-config)#aaa virtual-router 1 vrXyz1
host1(vr-group-config)#aaa virtual-router 2 vrXyz2
host1(vr-group-config)#aaa virtual-router 3 vrXyz3
host1(vr-group-config)#exit
host1(config)#
```

3. Enable broadcast accounting. Enter the correct virtual router context, and specify the virtual router group whose virtual routers will receive the accounting information.

```
host1(config)#virtual-router opVr100
host1:opVr100(config)#aaa accounting broadcast groupXyzCompany
```

Related Documentation

- *aaa accounting broadcast*
- *aaa accounting vr-group*
- *virtual-router*

Overriding AAA Accounting NAS Information

AAA accounting packets normally include two RADIUS attributes—NAS-IP-Address [4] and NAS-Identifier [32]—of the virtual router that generates the accounting information. You can override the default configuration and specify that accounting packets from particular broadcast virtual routers instead include the NAS-IP-Address and NAS-Identifier attributes of the authenticating virtual router.

To override the normal AAA accounting NAS information, access the correct virtual router context, and use the **radius override nas-info** command. For example:

```
host1(config)#virtual-router vrXyz1
host1:vrXyz1(config)#radius override nas-info
host1:vrXyz1(config)#virtual-router vrXyz2
host1:vrXyz2(config)#radius override nas-info
host1:vrXyz3(config)#exit
host1(config)#
```

- Related Documentation**
- *radius override nas-info*
 - *virtual-router*

Collecting Accounting Statistics

You can use the **aaa accounting statistics** command to specify how the AAA server collects statistics on the sessions it manages. Use the **volume-time** keyword to specify that AAA notifies applications to collect a full set of statistics from each of their connections. Use the **time** keyword to specify that only the uptime status is collected for each connection. Collecting only uptime information reduces the amount of data sent to AAA and is a more efficient use of system resources for customers that do not need a full set of statistics. The router collects a full set of statistics by default.

- Related Documentation**
- *aaa accounting statistics*

Configuring RADIUS AAA Servers

The number of RADIUS servers you can configure depends on available memory. The router has an embedded RADIUS client for authentication and accounting.



NOTE: You can configure B-RAS with RADIUS accounting, but without RADIUS authentication. In this configuration, the username and password on the remote end are not authenticated and can be set to any value.

You must assign an IP address to a RADIUS authentication or accounting server to configure it.

If you do not configure a primary authentication or accounting server, all authentication and accounting requests will fail. You can configure other servers as backup in the event that the primary server cannot be reached. Configure each server individually.

To configure an authentication or accounting RADIUS server:

1. Specify the authentication or accounting server address.

```
host1(config)#radius authentication server 10.10.10.1
host1(config-radius)#
or
host1(config)#radius accounting server 10.10.10.6
```

host1(config-radius)#

2. (Optional) Specify a UDP port for RADIUS authentication or accounting server requests.

host1(config-radius)#udp-port 1645

3. Specify an authentication or accounting server secret.

host1(config-radius)#key gismo

4. (Optional) Specify the number of retries the router makes to an authentication or accounting server before it attempts to contact another server.

host1(config-radius)#retransmit 2

5. (Optional) Specify the number of seconds between retries.

host1(config-radius)#timeout 5

6. (Optional) Specify the maximum number of outstanding requests.

host1(config-radius)#max-sessions 100

7. (Optional) Specify the amount of time to remove a server from the available list when a timeout occurs.

host1(config-radius)#deadtime 10

8. (Optional) In Global Configuration mode, specify whether the E Series router should move on to the next RADIUS server when the router receives an Access-Reject message for the user it is authenticating.

host1(config)#radius rollover-on-reject enable

9. (Optional) Enable duplicate address checking.

host1(config)#aaa duplicate-address-check enable

10. (Optional) Specify that duplicate accounting records be sent to the accounting server for a virtual router.

host1(config)#aaa accounting duplication routerBoston

11. (Optional) Enter the correct virtual router context, and specify the virtual router group to which broadcast accounting records are sent.

host1(config)#virtual-router vrSouth25

host1:vrSouth25(config)#aaa accounting broadcast westVrGroup38

host1:vrSouth25(config)#exit

12. (Optional) Specify that immediate accounting updates be sent to the accounting server when a response is received to an Acct-Start message.

host1(config)#aaa accounting immediate-update

13. (Optional) Specify whether the router collects all statistics or only the uptime status.

host1(config)#aaa accounting time

14. (Optional) Specify that tunnel accounting be enabled or disabled.

host1(config)#radius tunnel-accounting enable

15. (Optional) Specify the default authentication and accounting methods for the subscribers.

```
host1(config)#aaa authentication ppp default radius none
```

16. (Optional) Disable UDP checksums on virtual routers you configure for B-RAS.

```
host1:(config)#virtual router boston
host1:boston(config)#radius udp-checksum disable
```

Related Documentation

- *aaa accounting broadcast*
- *aaa accounting duplication*
- *aaa accounting immediate-update*
- *aaa authentication default*
- *aaa duplicate-address-check*
- *key*
- *max-sessions*
- *radius accounting server*
- *radius authentication server*
- *radius rollover-on-reject*
- *radius tunnel-accounting*
- *radius udp-checksum*
- *retransmit*
- *timeout*
- *udp-port*
- *virtual-router*

Configuring SNMP Traps

This section describes how to configure the router to send traps to SNMP when RADIUS servers fail to respond to messages, and how to configure SNMP to receive the traps.

To set up the router to send traps:

1. (Optional) Enable SNMP traps when a particular RADIUS authentication server fails to respond to Access-Request messages.

```
host1(config)#radius trap auth-server-not-responding enable
```

2. (Optional) Enable SNMP traps when all of the configured RADIUS authentication servers on a VR fail to respond to Access-Request messages.

```
host1(config)#radius trap no-auth-server-responding enable
```

3. (Optional) Enable SNMP traps when a RADIUS authentication server returns to active service.

```
host1(config)#radius trap auth-server-responding enable
```

4. (Optional) Enable SNMP traps when a RADIUS accounting server fails to respond to a RADIUS accounting request.

host1(config)#radius trap acct-server-not-responding enable

5. (Optional) Enable SNMP traps when all of the RADIUS accounting servers on a VR fail to respond to a RADIUS accounting request.

host1(config)#radius trap no-acct-server-responding enable

6. (Optional) Enable SNMP traps when a RADIUS accounting server returns to active service.

host1(config)#radius trap acct-server-responding enable

To set up SNMP to receive RADIUS traps:

1. Set up the appropriate SNMP community strings.

host1(config)#snmp-server community admin view everything rw

host1(config)#snmp-server community private view user rw

host1(config)#snmp-server community public view everything ro

2. Specify the interface whose IP address is the source address for SNMP traps.

host1(config)#snmp-server trap-source fastEthernet 0/0

3. Configure the host that should receive the SNMP traps.

host1(config)#snmp-server host 10.10.132.93 version 2c 3 udp-port 162 radius

4. Enable the SNMP router agent to receive and forward RADIUS traps.

host1(config)#snmp-server enable traps radius

5. Enable the SNMP on the router.

host1(config)#snmp-server



NOTE: For more information about these SNMP commands, see *JunosE System Basics Configuration Guide*.

Related Documentation

- *radius trap acct-server-responding*
- *radius trap acct-server-not-responding*
- *radius trap no-acct-server-responding*
- *radius trap auth-server-responding*
- *radius trap auth-server-not-responding*
- *radius trap no-auth-server-responding*
- *snmp-server*
- *snmp-server community*
- *snmp-server enable traps*
- *snmp-server host*

- *snmp-server trap-source*

Creating the AAA Local Authentication Environment

To create your local authentication environment:

1. Create local user databases—Create the default database or a named database.
2. Add entries to local user databases—Add user entries to the database. A database can contain information for multiple users.
3. Assign a local user database to the virtual router—Specify the database that the virtual router will use to authenticate subscribers.
4. Enable local authentication on the virtual router—Specify the **local** method as an AAA authentication method used by the virtual router.

Related Documentation

- [Creating AAA Local User Databases on page 73](#)
- [Adding AAA User Entries to Local User Databases on page 74](#)
- [Adding AAA User Entries to Default Local User Databases on page 74](#)
- [Configuring AAA User Entries in Local User Databases on page 75](#)
- [Assigning a Local User Database to a Virtual Router on page 76](#)
- [Enabling Local Authentication on the Virtual Router on page 77](#)

Creating AAA Local User Databases

When a subscriber connects to an E Series router that is using local authentication, the local authentication server uses the entries in the local user database selected by the virtual router to authenticate the subscriber.

A local authentication server can have multiple local user databases, and each database can have entries for multiple subscribers. The default local user database, if it exists, is used for local authentication by default. The E Series router supports a maximum of 100 user entries. A maximum of 100 databases can be configured.

To create a local user database, use the **aaa local database** command and the name of the database; use the name **default** to create the default local user database:

```
host1(config)#aaa local database westLocal40
```

Related Documentation

- [Adding AAA User Entries to Local User Databases on page 74](#)
- [Adding AAA User Entries to Default Local User Databases on page 74](#)
- [Configuring AAA User Entries in Local User Databases on page 75](#)
- [Assigning a Local User Database to a Virtual Router on page 76](#)
- [Enabling Local Authentication on the Virtual Router on page 77](#)

- *aaa local database*

Adding AAA User Entries to Local User Databases

The local authentication server uses the information in a local user database to authenticate a subscriber. A local user database can contain information for multiple users.

The E Series router provides two commands for adding entries to local user databases: the **username** command and the **aaa local username** command. You can specify the following parameters:

- Username—Name associated with the subscriber.
- Passwords and secrets—Single words that can be encrypted or unencrypted. Passwords use two-way encryption, and secrets use one-way encryption. Both passwords and secrets can be used with PAP authentication; however, only passwords can be used with CHAP authentication.
- IP address—The IP address to assign to the subscriber (**aaa local username** command only).
- IP address pool—The IP address pool used to assign the subscriber's IP address (**aaa local username** command only).
- Operational virtual router—The virtual router to which the subscriber is assigned. This parameter is applicable only if the subscriber is authenticated by the default virtual router (**aaa local username** command only).

Related Documentation

- [Creating AAA Local User Databases on page 73](#)
- [Adding AAA User Entries to Default Local User Databases on page 74](#)
- [Configuring AAA User Entries in Local User Databases on page 75](#)
- [Assigning a Local User Database to a Virtual Router on page 76](#)
- [Enabling Local Authentication on the Virtual Router on page 77](#)
- *aaa local username*
- *username*

Adding AAA User Entries to Default Local User Databases

The **username** command is similar to the command used by some third-party vendors. The command can be used to add entries in the default local user database; it is not supported for named local user databases. The IP address, IP address pool, and operational virtual router parameters are not supported in the **username** command. However, after the user is added to the default local user database, you can use the **aaa local username** command with a database name **default** to enter Local User Configuration mode and add the additional parameters.



NOTE: If the default local user database does not exist, the **username** command creates this database and adds the user entry to the database.

To add a subscriber and password or secret to the default local user database, complete the following step:

```
host1(config)#username rockyB password rockyPassword
```

Related Documentation

- [Creating AAA Local User Databases on page 73](#)
- [Adding AAA User Entries to Local User Databases on page 74](#)
- [Configuring AAA User Entries in Local User Databases on page 75](#)
- [Assigning a Local User Database to a Virtual Router on page 76](#)
- [Enabling Local Authentication on the Virtual Router on page 77](#)
- *username*

Configuring AAA User Entries in Local User Databases

To enter Local User Configuration mode and add user entries to a local user database, use the following commands:

1. Specify the subscriber's username and the database you want to use. Use the database name **default** to specify the default local user database. This command also puts the router into Local User Configuration mode.

```
host1(config)# aaa local username cksmith database westLocal40
host1(config-local-user)#
```



NOTE: You can use the **aaa local username** command to add or modify user entries to a default database that was created by the **username** command.

2. (Optional) Specify the type of encryption algorithm and the password or secret that the subscriber must use to connect to the router. A subscriber can be assigned either a password or a secret, but not both. For example:

```
host1(config-local-user)#password 8 iTtakes2%
```

3. (Optional) Specify the IP address to assign to the subscriber.

```
host1(config-local-user)#ip-address 192.168.101.19
```

4. (Optional) Specify the IP address pool used to assign the subscriber's IP address.

```
host1(config-local-user)#ip-address-pool svPool2
```

5. (Optional) Assign the subscriber to an operational virtual router. This parameter is applicable only if the subscriber is authenticated in the default virtual router.

```
host1(config-local-user)#operational-virtual-router boston2
```

Related Documentation

- [Creating AAA Local User Databases on page 73](#)
- [Adding AAA User Entries to Local User Databases on page 74](#)
- [Adding AAA User Entries to Default Local User Databases on page 74](#)
- [Assigning a Local User Database to a Virtual Router on page 76](#)
- [Enabling Local Authentication on the Virtual Router on page 77](#)
- *aaa local username*
- *ip-address*
- *ip-address-pool*
- *operational-virtual-router*
- *password*

Assigning a Local User Database to a Virtual Router

Use the procedure in this section to assign a local user database to a virtual router. The virtual router uses the database for local authentication when the subscriber connects to the E Series router. Use the following commands in Global Configuration mode:



NOTE: If you do not specify a local user database, the virtual router selects the default database by default. This applies to all virtual routers.

1. Specify the virtual router name.
`host1(config)# virtual-router cleveland`
2. Specify the database to use for authentication on this virtual router.
`host1:cleveland(config)# aaa local select database westLocal40`

Related Documentation

- [Creating AAA Local User Databases on page 73](#)
- [Adding AAA User Entries to Local User Databases on page 74](#)
- [Adding AAA User Entries to Default Local User Databases on page 74](#)
- [Configuring AAA User Entries in Local User Databases on page 75](#)
- [Enabling Local Authentication on the Virtual Router on page 77](#)
- *aaa local select database*
- *virtual-router*

Enabling Local Authentication on the Virtual Router

On the E Series router, RADIUS is the default AAA authentication method for PPP subscribers. Use the commands in this section to specify that the local authentication method is used.

To enable local authentication on the default router, use the following command:

```
host1(config)# aaa authentication ppp default local
```

To enable local authentication on a specific virtual router, first select the virtual router:

```
host1(config)# virtual-router cleveland
host1:cleveland(config)# aaa authentication ppp default local
```

Related Documentation

- [Creating AAA Local User Databases on page 73](#)
- [Adding AAA User Entries to Local User Databases on page 74](#)
- [Adding AAA User Entries to Default Local User Databases on page 74](#)
- [Configuring AAA User Entries in Local User Databases on page 75](#)
- [Assigning a Local User Database to a Virtual Router on page 76](#)
- *aaa authentication default*
- *virtual-router*

Example: Configuring AAA Local Authentication

This example creates a sample local authentication environment. The steps in this example:

1. Create a named local user database (**westfordLocal40**).
2. Configure the database **westfordLocal40**.
 - Add users **btjones** and **maryrdavis** and their attributes to the database.
3. Create the default local database using the optional **username** command.
 - Add optional subscriber parameters for user **cksmith** to the default database.
4. Assign the default local user database to virtual router **cleveland**; assign database **westfordLocal40** to the default virtual router and to virtual router **chicago**.
5. Enable AAA authentication methods **local** and **none** on all virtual routers.
6. Use the **show** commands to display information for the local authentication environment (various **show** command displays are listed after the example).

Example 1 This example shows the commands you use to create the AAA local authentication environment.

```
host1(config)#aaa local database westfordLocal40
```

```
host1(config)#aaa local username btjones database westfordLocal40
host1(config-local-user)#secret 38schillCy
host1(config-local-user)#ip-address-pool addressPoolA
host1(config-local-user)#operational-virtual-router boston2
host1(config-local-user)#exit
host1(config)#aaa local username maryrdavis database westfordLocal40
host1(config-local-user)#secret 0 davisSecret99
host1(config-local-user)#ip-address 192.168.20.106
host1(config-local-user)#operational-virtual-router boston1
host1(config-local-user)#exit
host1(config)#username cksmith password 0 yourPassword1
host1(config)#aaa local username cksmith database default
host1(config-local-user)#ip-address-pool addressPoolA
host1(config-local-user)#operational-virtual-router boston2
host1(config-local-user)#exit
host1(config)#virtual-router cleveland
host1(config)#aaa local select database default
host1(config)#virtual-router default
host1(config)#aaa local select database westfordLocal40
host1(config)#virtual-router chicago
host1(config)#aaa local select database westfordLocal40
host1(config)#virtual-router default
host1(config)#aaa authentication ppp default local none
```

Example 2 This example verifies that local authentication is configured on the router.

```
host1#show aaa authentication ppp default
local none
```

Example 3 This example uses the **show configuration category aaa local-authentication** command with the **databases** keyword to show the local user databases that are configured on the router.

```
host1# show configuration category aaa local-authentication databases
! Configuration script being generated on TUE NOV 09 2004 12:50:18 UTC
! Juniper Edge Routing Switch ERX1400
! Version: 6.1.0 (November 8, 2004 18:31)
! Copyright (c) 1999-2004 Juniper Networks, Inc. All rights reserved.
!
! Commands displayed are limited to those available at privilege level 15
!
! NOTE: This script represents only a subset of the full system configuration.
! The category displayed is: aaa local-authentication databases
!
hostname host1
aaa new-model
aaa local database default
aaa local database westfordLocal40
```

Example 4 This example uses the **local-authentication users** keywords to show the configured users and their parameters. The password for **username cksmith** is displayed unencrypted because the default setting of disabled or no for the **service password-encryption** command is used for the example. Secrets are always displayed encrypted.

```
host1# show configuration category aaa local-authentication users
! Configuration script being generated on THU NOV 11 2004 13:40:41 UTC
! Juniper Edge Routing Switch ERX1400
```

```

! Version: 6.1.0 (November 10, 2004 21:15)
! Copyright (c) 1999-2004 Juniper Networks, Inc. All rights reserved.
!
! Commands displayed are limited to those available at privilege level 15
!
! NOTE: This script represents only a subset of the full system configuration.
! The category displayed is: aaa local-authentication users
!
hostname host1
aaa new-model
aaa local username cksmith database default
    password yourPassword1
    operational-virtual-router boston2
    ip-address-pool addressPoolA
!
aaa local username btjones database westfordLocal40
    secret 5 }9s7-4N<WK2)2=)^!6~#
    operational-virtual-router boston2
    ip-address-pool addressPoolA
!
aaa local username maryrdavis database westfordLocal40
    secret 5 E@A:nDXJJ<irb\`mF#[j
    operational-virtual-router boston1
    ip-address 192.168.20.106

```

Example 5 This example uses the **users include-defaults** keywords to show the configured users and their parameters, including the default parameters **no ip-address** and **no ip-address-pool**.

```

host1# show configuration category aaa local-authentication users include-defaults
! Configuration script being generated on TUE NOV 09 2004 13:09:03 UTC
! Juniper Edge Routing Switch ERX1400
! Version: 6.1.0 (November 8, 2004 18:31)
! Copyright (c) 1999-2004 Juniper Networks, Inc. All rights reserved.
!
! Commands displayed are limited to those available at privilege level 15
!
! NOTE: This script represents only a subset of the full system configuration.
! The category displayed is: aaa local-authentication users
!
hostname host1
aaa new-model
aaa local username cksmith database default
    password yourPassword1
    operational-virtual-router boston2
    no ip-address
    ip-address-pool addressPoolA
!
aaa local username btjones database westfordLocal40
    secret 5 }9s7-4N<WK2)2=)^!6~#
    operational-virtual-router boston2
    no ip-address
    ip-address-pool addressPoolA
!
aaa local username maryrdavis database westfordLocal40
    secret 5 E@A:nDXJJ<irb\`mF#[j
    operational-virtual-router boston1
    ip-address 192.168.20.106
    no ip-address-pool

```

Example 6 This example uses the **virtual-router** keyword with the **default** specification to show the local user database that is used by the default virtual router.

```
host1# show configuration category aaa local-authentication virtual-router default
! Configuration script being generated on TUE NOV 09 2004 13:09:45 UTC
! Juniper Edge Routing Switch ERX1400
! Version: 6.1.0 (November 8, 2004 18:31)
! Copyright (c) 1999-2004 Juniper Networks, Inc. All rights reserved.
!
! Commands displayed are limited to those available at privilege level 15
!
! NOTE: This script represents only a subset of the full system configuration.
! The category displayed is: aaa local-authentication
!
virtual-router default
aaa local select database westfordLocal140
```

Example 7 This example uses the **virtual-router** keyword with a named virtual router. The **include-defaults** keyword shows the default configuration, including the line showing that there is no named local user database selected.

```
host1# show configuration category aaa local-authentication virtual-router cleveland
include-defaults
! Configuration script being generated on TUE NOV 09 2004 13:09:25 UTC
! Juniper Edge Routing Switch ERX1400
! Version: 6.1.0 (November 8, 2004 18:31)
! Copyright (c) 1999-2004 Juniper Networks, Inc. All rights reserved.
!
! Commands displayed are limited to those available at privilege level 15
!
! NOTE: This script represents only a subset of the full system configuration.
! The category displayed is: aaa local-authentication
!
virtual-router cleveland
no aaa local select
```

**Related
Documentation**

- *aaa authentication default*
- *aaa local database*
- *aaa local select database*
- *aaa local username*
- *ip-address*
- *ip-address-pool*
- *operational-virtual-router*
- *password*
- *secret*
- *show aaa authentication default*
- *show configuration*
- *virtual-router*

Configuring DNS Primary and Secondary NMS

To configure the DNS primary and secondary name server addresses:

1. Specify the IP address of the DNS primary name server.

```
host1(config)#aaa dns primary 10.10.10.5
```

or, for IPv6,

```
host1(config)#aaa ipv6-dns primary 2001:db8::8001
```

2. Specify the IP address of the DNS secondary name server.

```
host1(config)#aaa dns secondary 10.10.10.6
```

or, for IPv6,

```
host1(config)#aaa ipv6-dns secondary 2001:db8::8002
```



NOTE: The router uses name server addresses exclusively for PPP clients and not for domain name server resolution.

Related Documentation

- *aaa dns*
- *aaa ipv6-dns*

Configuring WINS Primary and Secondary NMS

To configure the WINS primary and secondary name server addresses:

1. Specify the IP address of the WINS primary name server.

```
host1(config)#aaa wins primary 192.168.10.05
```

2. Specify the IP address of the WINS secondary name server.

```
host1(config)#aaa wins secondary 192.168.10.40
```



NOTE: The router uses name server addresses exclusively for PPP clients and not for domain name server resolution.

Related Documentation

- *aaa wins*

Configuring a Local Address Server

You can create, modify, and delete address pools. You can display address pool information or status with the **show ip local pool** command. The following are examples of tasks you can configure:

- Specify an addressing scheme.

```
host1(config)#ip address-pool local
```

- Map an address pool name to a range of local addresses. You can also use this command to add additional ranges to a pool.

```
host1(config)#ip local pool addrpool_10 192.168.56.10 192.168.56.15
```

- Map a primary local address pool name to a domain name.

```
host1(config)#aaa domain-map westford.com
host1(config-domain-map)#address-pool-name poolA
```

- (Optional) Map a backup address pool to a domain name, which is used for address allocation if the primary local address pool is fully allocated.

```
host1(config)#aaa domain-map westford.com
host1(config-domain-map)#backup-address-pool-name backup_poolB
```

- (Optional) Map the domain name to the IPv6 local address pool, which is used for prefix delegation. If the authentication server returns the prefix pool name in the Framed-Ipv6-Pool attribute of the RADIUS-Accept-Request message, this value overrides the IPv6 local pool configured using the **ipv6-prefix-pool-name** command.

```
host1(config)#aaa domain-map westford.com
host1(config-domain-map)#ipv6-prefix-pool-name local_addr_pool
```

- Delete an address pool.

```
host1(config)#no ip local pool addrpool_10
```



NOTE: If a pool or range is deleted and addresses are outstanding, the AAA server logs out the clients using the addresses.

- Create a shared local address pool.

```
host1(config)#ip local shared-pool Shared_LAS_Pool_A DHCP_Pool_1
```

- Delete a shared local address pool.

```
host1(config)#no ip local shared-pool Shared_LAS_Pool_C
```

- Set SNMP variables by specifying an existing pool name and values.

```
host1(config)#ip local pool addrpool_10 warning 90 80
```

Related Documentation

- *aaa domain-map*
- *address-pool-name*
- *backup-address-pool-name*
- *ip address-pool*
- *ip local pool*
- *ip local shared-pool*
- *ipv6-prefix-pool-name*

Creating an IP Interface


You can configure IP interfaces that support the following configurations:

- [Configuring Single PPP Clients per ATM Subinterface on page 83](#)
- [Configuring Multiple PPP Clients per ATM Subinterface on page 84](#)

Configuring Single PPP Clients per ATM Subinterface

Figure 3 on page 83 shows a conceptual view of the configuration of a single PPP client per ATM subinterface.

Figure 3: Single PPP Clients per ATM Subinterface



Configure an ATM interface by entering Configuration mode and performing the following tasks. For more information about configuring ATM interfaces, see *JunosE Link Layer Configuration Guide*.

1. Configure a physical interface.

```
host1(config)#interface atm 0/1
```

2. Configure the subinterface.

```
host1(config-if)#interface atm 0/1.20
```

3. Configure a permanent virtual circuit (PVC) by specifying the vcd (virtual circuit descriptor), the vci (virtual channel identifier), the vpi (virtual path identifier), and the encapsulation type.

```
host1(config-if)#atm pvc 10 22 100 aal5snap
```

4. Configure PPP encapsulation.

```
host1(config-if)#encapsulation ppp
```

5. Configure PAP or CHAP authentication.

```
host1((config-if))#ppp authentication chap
```

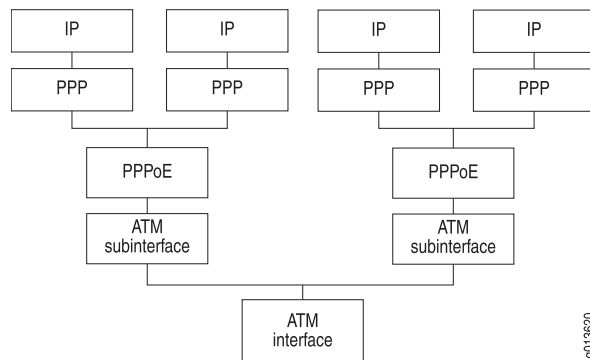
6. Assign a profile to the PPP interface.

```
host1(config-subif)#profile foo
```

Configuring Multiple PPP Clients per ATM Subinterface

Figure 4 on page 84 shows how PPPoE supports multiplexing of multiple PPP sessions per ATM subinterface.

Figure 4: Multiple PPP Clients per ATM Subinterface



9013620

Configure an ATM interface by entering Configuration mode and performing the following tasks. For more information about configuring ATM interfaces, see *JunosE Link Layer Configuration Guide*.

1. Configure a physical interface.

```
host1(config)#interface atm 0/1
```

2. Configure the subinterface.

```
host1(config-if)#interface atm 0/1.20
```

3. Configure a PVC by specifying the vcd (virtual circuit descriptor), the vci (virtual channel identifier), the vpi (virtual path identifier), and the encapsulation type.

```
host1(config-if)#atm pvc 10 22 100 aal5snap
```

4. Configure PPPoE encapsulation.

```
host1(config-if)#encapsulation pppoe
```

5. Configure the subinterface for one PPP client.

```
host1(config-if)#interface atm 0/1.20.1
```

6. Configure PPP encapsulation.

```
host1(config-if)#encapsulation ppp
```

7. Configure PAP or CHAP authentication.

```
host1((config-if))#ppp authentication chap
```

8. Apply the profile to the PPP interface.

```
host1(config-subif)#profile foo2
```

9. Configure the subinterface for a second PPP client.

```
host1(config-if)#interface atm 0/1.20.2
```

10. Configure PPP encapsulation.

```
host1(config-if)#encapsulation ppp
```

11. Configure PAP or CHAP authentication.

```
host1((config-if))#ppp authentication chap
```

12. Apply the profile to the PPP interface.

```
host1(config-subif)#profile foo2
```

Related Documentation

- *atm pvc*
- *encapsulation ppp*
- *interface*
- *ppp authentication*
- *profile*

Controlling Access to Domain Names

You can control a PPP subscriber's access to certain domains on given interfaces. As the administrator, you can use the **deny** command to prevent PPP subscribers from using unauthorized domain names. Using the **allow** command, you can allow PPP subscribers to use authorized domain names.

In this example, the administrator wants to restrict access of a PPP interface to the specific domain **abc.com**.

1. Create an AAA profile.

```
host1(config)#aaa profile restrictToABC
```

2. Specify the domain name you want to allow.

```
host1(config-aaa-profile)#allow abc.com
```

3. Specify the domain name you want to restrict.

```
host1(config-aaa-profile)#deny default
```

4. Associate the AAA profile to the designated PPP interface.

```
host1(config-if)#ppp aaa-profile restrictToABC
```

When configured as such, the following is a likely scenario:

- PPP passes the AAA profile **restrictToABC** to AAA in the authentication request.
- AAA performs the following:
 - Receives the authentication request from PPP with the subscriber's name **will@xyz.com**.
 - Parses the domain name **xyz.com** and examines the specified AAA profile **restrictToABC**.
 - Determines that the AAA profile **restrictToABC** is valid.

- Searches **restrictToABC** for a match on the PPP subscriber's domain name and finds no match.
- Searches **restrictToABC** for a match on the domain name **default**.
- Finds a match and denies the user access.

**Related
Documentation**

- *aaa profile*
- *allow*
- *deny*
- *ppp aaa-profile*

Example: Associating all Subscribers of a PPP Interface with a Specific Domain Name

In this example, an administrator wants to associate all subscribers of a PPP interface with a specific domain name.

1. Create an AAA profile.

```
host1(config)#aaa profile forwardToXyz
```

2. Map the original domain name to the mapped domain name for domain map lookup.

```
host1(config-aaa-profile)#translate default xyz.com
```

3. Associate the AAA profile with the designated PPP interface.

```
host1(config-if)#ppp aaa-profile forwardToXyz
```

When configured as such, the following scenario is typical:

- PPP passes the AAA profile **forwardToXyz** to AAA in the authentication request.
- AAA performs the following tasks:
 - Receives the authentication request from PPP with the subscriber's name **morris@abc.com**.
 - Parses the domain name **abc.com** and examines the specified AAA profile **forwardToXyz**.
 - Determines that the AAA profile **forwardToXyz** is valid.
 - Searches **forwardToXyz** for a match on the PPP subscriber's domain name and finds no match.
 - Searches **forwardToXyz** for a match on the domain name **default**.
 - Finds a match and continues as normal using the domain name **xyz.com**.



NOTE: If there is no matching entry in the AAA profile for the user's domain name or for the domain name **default**, then AAA continues processing as if there were no AAA profile.

If the user's name does not contain a domain name, then AAA attempts to match to the domain name **none** in the AAA profile. If there is no entry for **none**, then AAA attempts to match for the domain name **default** in the AAA profile. If there is no entry for either **none** or **default**, then AAA continues processing as if there were no AAA profile.

Related Documentation

- *aaa profile*
- *allow*
- *deny*
- *ppp aaa-profile*
- *translate*

Example: Associating Multiple Domain Names with a Specific Domain Name

In this example, an administrator wants to use aliases; that is, to associate multiple domain names with a specific domain name and not allow other domain names.

1. Create an AAA profile.

```
host1(config)#aaa profile toAbc
```

2. Map the original domain name to the mapped domain name for domain map lookup.

```
host1(config-aaa-profile)#translate abc1.com abc.com
host1(config-aaa-profile)#translate abc2.com abc.com
host1(config-aaa-profile)#translate abc3.com abc.com
```

3. Specify the domain name you want to restrict.

```
host1(config-aaa-profile)#deny default
```

4. Associate the AAA profile with the designated PPP interface.

```
host1(config-if)#ppp aaa-profile toAbc
```

When configured as such, the following scenario is typical:

- PPP passes the AAA profile toAbc to AAA in the authentication request.
- AAA:
 - Receives the authentication request from PPP with the subscriber's name **jane@abc1.com**
 - Parses the domain name **abc1.com** and examines the specified AAA profile toAbc
 - Determines that the AAA profile **toAbc** is valid

- Searches **toAbc** for a match on the PPP subscriber's domain name and finds a match
- Continues as normal using the domain name **abc.com**



NOTE: If there is no matching entry in the AAA profile for the user's domain name or for the domain name **default**, then AAA continues processing as if there were no AAA profile.

If the user's name does not contain a domain name, then AAA attempts to match to the domain name **none** in the AAA profile. If there is no entry for **none**, then AAA attempts to match for the domain name **default** in the AAA profile. If there is no entry for either **none** or **default**, then AAA continues processing as if there were no AAA profile.

**Related
Documentation**

- *aaa profile*
- *allow*
- *deny*
- *ppp aaa-profile*
- *translate*

Configuring an AAA Per-Profile Attribute List

JunosE Software enables you to configure AAA-specific attributes for subscribers attached to a specific PPP profile. If a per-profile list is configured, then only the attributes specified in the per-profile list are processed. If the per-profile list is not configured, then the existing standard attributes are configured.



NOTE: The attributes supported by the per-profile list take precedence over the standard AAA attribute configuration. By default, the inclusion of all attributes is disabled in the per-profile list.

This feature enables you to configure the following AAA attributes:

- **tunnel ignore nas-port**
- **tunnel ignore nas-port-type**

In this example, AAA-specific attributes are configured for subscribers attached to a specific PPP profile. You can configure this as follows:

1. Create an AAA per-profile attribute list, and configure the required AAA attributes in the list.

```
host1(config)#aaa per-profile-attr-list abc
host1 (config-perprofile-list)#action-type enable
```



```
host1 (config-perprofile-list)#attributes tunnel-ignore-nasport
tunnel-ignore-nasport-type
```

2. Create an AAA profile.

```
host1(config)#aaa profile aaaprofile1
```

3. Specify the AAA attribute list in the AAA profile.

```
host1(config-aaa-profile)#aaa-perprofilelist-name abc
```

4. Create a PPP profile.

```
host1(config)#profile pppprofile1
```

5. Attach the AAA profile name to the PPP profile.

```
host1(config-profile)#ppp aaa-profile aaaprofile1
```

6. To view the attributes configured in the AAA per-profile attribute list, issue the **show aaa per-profile-attr-list** command.

```
host1#show aaa per-profile-attr-list abc
Profile name: abc
Attribute Name          Status
-----
tunnel-ignore-nasport   enabled
tunnel-ignore-nasport-type enabled
```

Related Documentation

- *aaa profile*
- *aaa-perprofilelist-name*
- *aaa per-profile-attr-list (For Global Configuration)*
- *action-type*
- *attributes (AAA)*
- *ppp aaa-profile*
- *profile*
- *show aaa per-profile-attr-list*

Configuring Interim, Broadcast, and Policy-Based Accounting in Virtual Router Groups and PPP Profiles

You can use the **aaa virtual-router** command with the **interim-update** or **policy-accounting** keyword to enable or disable sending of periodic interim accounting messages or policy-based accounting statistics to broadcast accounting servers configured in a broadcast virtual router group. The policy-based accounting messages sent from AAA to a RADIUS server contains the Calling-Station-Id [31] and Event-Timestamp [55] RADIUS attributes.

You can use the **aaa accounting interim-update** command to enable or disable periodic sending of interim accounting messages to a primary accounting server on a per-virtual router basis. The transmission interval of the interim accounting messages depends on the user accounting interval configured using the **aaa user accounting interval**.

When the user accounting interval is set as 0, interim accounting is not performed for users attached to the virtual router.

You can use the ***ppp aaa-accounting-broadcast*** command to assign a broadcast virtual router group to a PPP profile so as to enable broadcast accounting in PPP profiles.



NOTE:

- When a broadcast virtual router group is configured at both PPP profile and virtual router levels, AAA sends accounting messages only to broadcast accounting servers in the broadcast virtual router group that is configured at both PPP profile and virtual router levels.
- Policy-based accounting is not supported on a primary accounting server.
- IP version 6 (IPv6) policy-based accounting is not supported for the IPv6 policy returned from the RADIUS server.
- The configurations done using the ***aaa virtual-router***, ***aaa accounting interim-update***, and ***ppp aaa accounting broadcast*** commands are saved across the chassis cluster and unified in-service software upgrade (ISSU).

To configure interim, broadcast, and policy-based accounting in virtual router groups and PPP profiles:

- Enable sending of the Interim-Acct request periodically at a configured user accounting interval to the primary accounting server.

```
host1(config)#aaa accounting interim-update enable
```

- Create a virtual router group and enter VR Group Configuration mode.

```
host1(config)#aaa accounting vr-group groupxyz
host1(vr-group-config)#
```

- Add a virtual router to the group and enable sending of periodic interim accounting messages and policy-based accounting statistics to broadcast accounting servers configured in the virtual router group.

```
host1(vr-group-config)#aaa virtual-router 1 vxxyz1 interim-accounting enable
host1(vr-group-config)#aaa virtual-router 1 vxxyz1 policy-accounting enable
host1(vr-group-config)#exit
host1(config)#
```

- Enable broadcast accounting.

```
host1(config)#virtual-router vr4
host1:vr4(config)#aaa accounting broadcast groupxyz
host1:vr4(config)#virtual-router default
host1(config)#
```

- Create a PPP profile and assign the virtual router group to enable broadcast accounting in the PPP profile.

```
host1(config)#profile pppprofile10
host1(config-profile)#ppp aaa-accounting-broadcast groupxyz
```

- Related Documentation**
- [RADIUS Authentication and Accounting Servers Configuration Overview on page 17](#)
 - [Broadband Remote Access Support for PPP Overview](#)
 - [Profile Characteristics](#)
 - [Monitoring AAA-Specific Virtual Router Groups on page 118](#)
 - [Monitoring AAA Accounting Configuration on page 115](#)
 - `aaa accounting broadcast`
 - `aaa accounting vr-group`
 - `profile`
 - `virtual-router`

Configuring the NAS-Port-Type Attribute Manually

You can manually configure the NAS-Port-Type RADIUS attribute (attribute 61) in AAA profiles for ATM and Ethernet interfaces. Doing so allows AAA profiles to determine the NAS port type for a given connection.

To set the NAS-Port-Type attribute for ATM or Ethernet interfaces:

1. Create an AAA profile.
`host1(config)#aaa profile nasPortType`
2. (Optional) Set the NAS-Port-Type attribute for ATM interfaces.
`host1(config-aaa-profile)#nas-port-type atm wireless-80211`
3. (Optional) Set the NAS-Port-Type attribute for Ethernet interfaces.
`host1(config-aaa-profile)#nas-port-type ethernet wireless-cable`

- Related Documentation**
- `aaa profile`
 - `nas-port-type atm`
 - `nas-port-type ethernet`

Configuring a Service Description for the AAA Profile

You can specify a service description that will be associated with an AAA profile. The description can then be exported through RADIUS by the Service-Description attribute (RADIUS attribute 26-53) in AAA profiles.

To set the Service-Description attribute:

1. Create the AAA profile.
`host1(config)#aaa profile xyzCorpPro2`
2. Set the Service-Description attribute.

```
host1(config-aaa-profile)#service-description bos-xyzcorp
```

- Related Documentation**
- *aaa profile*
 - *service-description*

Configuring the Route-Download Server to Download Routes

When you configure the E Series router as a route-download server, you specify the RADIUS server that you want to download the routes to your router. You can also modify the route-download server's default configuration parameters, such as when to start the download process each day, how often to download routes, and how long to wait after a download error before retrying the process.

- To configure a RADIUS route-download server to download IPv4 routes:
 1. Specify the IP address and the key of the RADIUS server that you want to download routes.

```
host1(config)#radius route-download server 192.168.1.17
host1(config-radius)#key 35radsrv92
```

2. (Optional) Specify the UDP port used for RADIUS route-download server requests.

```
host1(config-radius)#udp-port 1812
host1(config-radius)#exit
host1(config)#
```

3. Enable the route-download feature and optionally modify default parameters as needed.

```
host1(config)#aaa route-download 1200 retry-interval 25 password Configured
synchronization 03:45:00
```

4. (Optional) Verify your route-download configuration:

```
host1(config)#exit
host1#show aaa route-download

AAA Route Downloader:      configured in virtual router default
Download Interval:        1200 minutes
Retry Interval:           25 minutes
Default Cost:             2
Default Tag:              0
Base User Name:           <HOSTNAME>
Password:                 Configured
Synchronization:          03:45:00

Status:                   downloading
Last Download Attempt:    TUE FEB 9 22:07:30 2007
Last Download Success:    <NEVER>
Last Regular Download:    not complete
Next Download Scheduled:  <DOWNLOAD ACTIVE>
Next Regular Download:    WED FEB 9 22:27:00 2007
```

- To configure a RADIUS route-download server to download IPv6 routes:
 1. Specify the IPv6 address and the key of the RADIUS server that you want to download routes.

```
host1(config)#radius route-download server 192.168.1.17
host1(config-radius)#key 35radsrv92
```

2. (Optional) Specify the UDP port used for RADIUS route-download server requests.

```
host1(config-radius)#udp-port 1812
host1(config-radius)#exit
host1(config)#exit
```

3. Enable the route-download feature and optionally modify default parameters as needed.

```
host1(config)#aaa route-download ipv6
```

4. (Optional) Verify your route-download configuration:

```
host1(config)#exit
host1#show aaa route-download ipv6

AAA Route Downloader:      configured in virtual router default
Download Interval:         720 minutes
Retry Interval:            10 minutes
Default Cost:              2
Default Tag:               0
Base User Name:            <HOSTNAME>
Password:                  <DEFAULT>
Synchronization:          <NOT SET>

Status:                    idle
Last Download Attempt:     TUE DEC 13 2011 00:05:43 UTC
Last Download Success:     TUE DEC 13 2011 00:05:43 UTC
Last Regular Download:     complete
Next Download Scheduled:   TUE DEC 13 2011 12:05:42 UTC
Next Regular Download:     TUE DEC 13 2011 12:05:42 UTC
```



NOTE: If optional parameters such as retry-interval, synchronization, tag, cost, and download interval are configured for either IPv4 or IPv6 route downloads, they are applied to both IPv4 and IPv6 route downloads. However, the username and password are configured separately for IPv4 and IPv6 routes.

Related Documentation

- [aaa route-download](#)
- [aaa route-download ipv6](#)
- [key](#)
- [radius route-download server](#)
- [show aaa route-download](#)
- [udp-port](#)

Configuring the Router to Obtain the LLID for a Subscriber

To configure the router to obtain the LLID for a subscriber:

1. Create an AAA profile that supports subscriber preauthentication.

```
host1(config)#aaa profile preAuthLlid
host1(config-aaa-profile)#pre-authenticate
host1(config-aaa-profile)#exit
```

2. Define a RADIUS preauthentication server.

```
host1(config)#radius pre-authentication server 10.10.10.1
host1(config-radius)#key abc123
host1(config-radius)#exit
```

3. Associate the AAA profile with the designated PPP interface.

```
host1(config)#interface atm 4/3.101
host1(config-subif)#ppp aaa-profile preAuthLlid
```

4. (Optional) Verify that preauthentication support is configured for the AAA profile.

```
host1(config-subif)#run show aaa profile name PreAuthLlid
preAuthLlid:
  atm nas-port-type: ADLSL-CAP
  ethernet nas-port-type: Cable
  profile-service-description: xyzService
  pre-authenticate
  allow xyz.com
  deny default
  translate xyz1.com abc.com
```

For information, see [“Setting Baselines for Remote Access” on page 113](#).

5. (Optional) Verify configuration of the RADIUS preauthentication server.

```
host1(config-subif)#run show radius pre-authentication servers
```

RADIUS Pre-Authentication Configuration						
IP Address	Udp Port	Retry Count	Timeout	Maximum Sessions	Dead Time	Secret
10.10.10.1	1812	3	3	255	0	radius

You can also display configuration information for preauthentication servers by using the **show radius servers** command. For information, see [“Setting Baselines for Remote Access” on page 113](#).

6. (Optional) Display statistics for the RADIUS preauthentication server.

To display preauthentication statistics, use the **show radius pre-authentication statistics** command. For information, see [“Setting Baselines for Remote Access” on page 113](#).

To display a count of preauthentication requests and responses, use the **show aaa statistics** command. For information, see [“Setting Baselines for Remote Access” on page 113](#).

- Related Documentation**
- *aaa profile*
 - *interface*
 - *key*

- `ppp aaa-profile`
- `pre-authenticate`
- `radius pre-authentication server`
- `show aaa profile`
- `show radius servers`

Troubleshooting Subscriber Preauthentication

Problem Description: You can configure the router to send traps to SNMP when a RADIUS preauthentication server fails to respond to messages. To do so, you use the same procedure and commands as you do to configure SNMP traps for a RADIUS authentication server.

Solution For example, to enable SNMP traps when a particular RADIUS preauthentication server fails to respond to Access-Request messages, use the **radius trap auth-server-not-responding enable** command.

Related Documentation

- [Configuring SNMP Traps on page 71](#)
- `radius trap auth-server-not-responding`

Configuring Custom Mappings for PPP Terminate Reasons

This example describes a sample configuration procedure that creates custom mappings for PPP terminate reasons.

1. Configure the router to include the Acct-Terminate-Cause attribute in RADIUS Acct-Off messages.

```
host1(config)#radius include acct-terminate-cause acct-off enable
```

2. (Optional) Display the current PPP terminate-cause mappings.

```
host1(config)# run show terminate-code ppp
```

Apps	Terminate Reason	Description	Radius Code
ppp	authenticate-authenticator-timeout	authenticate authenticator timeout	17
ppp	authenticate-challenge-timeout	authenticate challenge timeout	10
ppp	authenticate-chap-no-resources	authenticate chap no resources	10
ppp	authenticate-chap-peer-authenticator-timeout	authenticate chap peer authenticator timeout	17
ppp	authenticate-deny-by-peer	authenticate deny by peer	17
ppp	authenticate-inactivity-timeout	authenticate inactivity timeout	4
--More--			

3. (Optional) Display all PPP terminate reasons.

```

host1(config)# terminate-code ppp ?
authenticate-authenticator-timeout      Configure authenticate
                                         authenticator timeout
                                         translation
authenticate-challenge-timeout          Configure authenticate
                                         challenge timeout translation
authenticate-chap-no-resources           Configure authenticate chap no
                                         resources translation
authenticate-chap-peer-
authenticate-authenticator-timeout       Configure authenticate chap
                                         peer authenticator timeout
                                         translation
authenticate-deny-by-peer               Configure authenticate deny by
                                         peer translation
--More--

```

4. Configure your customized PPP terminate-cause to RADIUS Acct-Terminate-Cause code mappings.

```

host1(config)#terminate-code ppp authenticate-authenticator-timeout radius 3
host1(config)#terminate-code ppp authenticate-challenge-timeout radius 4

```

5. Verify the new terminate-cause mappings.

```

host1(config)#run show terminate-code ppp

```

Apps	Terminate Reason	Description	Radius Code
ppp	authenticate-authenticator-timeout	authenticate authenticator timeout	3
ppp	authenticate-challenge-timeout	authenticate challenge timeout	4
ppp	authenticate-chap-no-resources	authenticate chap no resources	10
ppp	authenticate-chap-peer-authenticator-timeout	authenticate chap peer authenticator timeout	17
ppp	authenticate-deny-by-peer	authenticate deny by peer	17
ppp	authenticate-inactivity-timeout	authenticate inactivity timeout	4
ppp	authenticate-max-requests	authenticate max requests	10

--More--

- Related Documentation**
- *radius include*
 - *show terminate-code*
 - *terminate-code*

Configuring Duplicate IPv6 Prefix Check

You can enable detection of duplicates of IPv6 Neighbor Discovery router advertisement prefixes and DHCPv6 delegated prefixes.

To enable detection of duplicate IPv6 prefixes:

From Global Configuration mode, enable the prefix-checking capability

```

host1(config)#aaa duplicate-prefix-check enable

```


- Related Documentation**
- [Duplicate IPv6 Prefix Check Overview on page 42](#)
 - *aaa duplicate-prefix-check*

Configuring Detection of Duplicate IPv6 Prefixes in the AAA User Profile Database

You can enable detection of duplicates of IPv6 Neighbor Discovery router advertisement prefixes and DHCPv6 delegated prefixes in the AAA user profile database.

To enable enhanced detection of duplicate IPv6 prefixes:

- From Global Configuration mode, enable the enhanced duplicate IPv6 prefix-checking capability.

```
host1(config)#aaa duplicate-prefix-check-extension enable
```

- Related Documentation**
- [Duplicate IPv6 Prefix Detection in the AAA User Profile Database Overview on page 42](#)
 - [Monitoring Duplicate IPv6 Prefixes in the AAA User Profile Database on page 155](#)
 - *aaa duplicate-prefix-check-extension*

Configuring the SRC Client

You can configure SRC clients on a per-virtual-router basis. To configure the SRC client:

1. Enable the SRC client and plug-ins such as IP, DHCP, and DHCPv6. With the CLI **sscc enable** command, you can specify BER-encoded information exchange for COPS-PR.

```
host1(config)#sscc enable cops-pr
```

2. Specify the IP addresses of up to three service activation engines (SAEs) (primary, secondary, and tertiary). You can optionally specify the port on which the SAEs listen for activity.

```
host1(config)#sscc primary address
host1(config)#sscc secondary address 192.168.12.1 port 3288
```

3. (Optional) Enable policy and QoS configuration support for IPv6 interfaces.

```
host1(config)#sscc protocol ipv6
```

4. (Optional) Enable policy and QoS configuration support for L2TP interfaces on an L2TP access concentrator (LAC).

```
host1(config)#sscc protocol lac
```

5. (Optional) Specify on which router the TCP/COPS connection is to be established.

```
host1(config)#sscc transportRouter chicago
```



NOTE: If a COPS connection is in the open state (displayed in the “The Connection State is” field in the output of the `show sscd info` command), the router that you configure on which the COPS connection is to be established by using the `sscd transportRouter name` command does not take effect.

6. (Optional) Specify a fixed source address for the TCP/COPS connection created for an SRC client session.

```
host1(config)#sscd sourceAddress 10.9.123.8
```

7. (Optional) Specify a fixed source interface for the TCP/COPS connection.

```
host1(config)#sscd sourceInterface atm 3/0
```

8. (Optional) Specify the delay period during which an SRC client waits for a response from the SAE.

```
host1(config)#sscd retryTimer 120
```

9. (Optional) Specify the time period until which the rebooted SRC client waits for the primary SRC server to establish a connection with the SRC client. If the connection is not established within the configured time period, the SRC client selects the SRC server randomly, which is the default behavior of the SRC client.

```
host1(config)#sscd connectivityTimer 5
```



NOTE: The connectivity timer and ping operation to the primary SRC server are initiated when the first uplink or downlink line module other than the ES2 4G or ES2 10G ADV LMs with ES2-S1 Service IOA comes online. The ping operation is tried every 3 seconds after the timer is started until the timer stops or expires. The timer is stopped after a successful ping response is received.

10. (Optional) Enable the user IP address mask to be sent to an SRC server in place of the interface IP address mask for a virtual router.

```
host1(config)#sscd option user-ip-mask-override
```

11. (Optional) Enable the calling station ID to be sent to an SRC server for a virtual router.

```
host1(config)#sscd option send-calling-station-id
```

You can configure a virtual router to send the default calling station ID or the overridden calling station ID to an SRC server irrespective of the RADIUS settings. If you want to enable the SRC client to send the Calling-Station-Id [31] RADIUS attribute to the COPS server only if this attribute is included in the RADIUS Access-Request, Acct-Start, or Acct- Stop message, you can use the **radius-default-value** attribute with the **sscd option send-calling-station-id** command.

```
host1(config)#sscd option send-calling-station-id radius-default-value
```

If you want to enable the SRC client to send the Calling-Station-Id [31] RADIUS attribute to the COPS server, regardless of whether this attribute is included in the

RADIUS Access-Request, Acct-Start, or Acct- Stop message, you can use the **radius-override-value** attribute with the **sscc option send-calling-station-id** command.

```
host1(config)#sscc option send-calling-station-id radius-override-value
```

You must configure either the **radius calling-station-format** command or the **radius override calling-station-id remote-circuit-id** before you enable the functionality to cause the calling station ID to be always sent to an SRC server for a virtual router, regardless of whether the ID is included or excluded from the Access-Request and Acct-Start messages.



NOTE: If you did not configure Calling-Station-Id attribute format using the **radius calling-station-format** command or did not configure the PPPoE remote circuit ID to be used in RADIUS messages instead of Calling-Station-Id using the **radius override calling-station-id** command, the Calling-Station-Id attribute is sent to the COPS server from the SRC client only if this attribute is contained in the RADIUS messages. In such a scenario, the attribute is not sent from the SRC client to the COPS server even if you configured the **sscc option send-calling-station-id radius-override-value** command.

12. (Optional) Enable the local QoS profile attachment information to be sent to an SRC server for a virtual router.

```
host1(config)#sscc option send-local-qos-profile-config
```

13. (Optional) Enable the LAC-side NAS-IP address information to be sent to an SRC server for a virtual router.

```
host1(config)#sscc option send-lac-nas-ip
```

14. (Optional) Enable the LAC-side NAS-Port information to be sent to an SRC server for a virtual router.

```
host1(config)#sscc option send-lac-nas-port
```

15. (Optional) Enable the SRC client to obtain updated line-rate parameters from ANCP and transmit them to the COPS server.

```
host1(config)#sscc update-policy-request enable
```

16. (Optional) Restart a COPS connection to, and resynchronize with, an SRC server.

```
host1#sscc restart
```

Related Documentation

- [Monitoring the SRC Client Connection Status on page 155](#)
- *sscc address*
- *sscc connectivityTimer*
- *sscc enable*
- *sscc option*
- *sscc protocol ipv6*

- *sscc protocol lac*
- *sscc restart*
- *sscc retryTimer*
- *sscc sourceAddress*
- *sscc sourceInterface*
- *sscc transportRouter*
- *sscc update-policy-request enable*

Configuring the Forwarding of COPS Requests to the SRC Server Based on DCM Profiles

You can configure the SRC client on an E Series router, which functions as the Common Open Policy Service (COPS) client, to send COPS messages to the SRC server or the COPS server based on the dynamic configuration manager (DCM) profile. For subscribers that use PPP links to establish sessions with the router or the SRC client and for which subscriber policies are managed by the SRC software, you can configure the setting in the PPP profiles to enable the SRC client to send COPS messages to the SRC server. This method of transmission of COPS request messages to the SRC server facilitates effective, optimal control of subscriber login events in the SRC software.

To configure a PPP profile with the setting to send COPS requests to the SRC server:

1. Create a PPP profile.

```
host1(config)#profile pppprofile1
```

2. Configure the transmission of COPS request messages to the SRC server for all subscribers that are assigned this PPP profile.

```
host1(config)#ip send-cops-request
```

By default, COPS messages are sent to the SRC server. You must configure at least one IP configuration parameter in the PPP profile to enable the default behavior of the command to be effective. This functionality is applicable in environments where PPP links between the customer premises equipment (CPE) and the provider edge (PE) device or the router are configured for IPv4 or IPv6 subscriber sessions, either as independent or combined sessions. Also, this capability is effective only for dynamic PPP subscribers and not for DHCP and static subscriber sessions.

Use the **no** version to disable the transmission of COPS messages from the SRC client to the SRC server for PPP subscribers.

Related Documentation

- *ip send-cops-request*

Configuring the DHCPv6 Local Address Pools

The IPv6 local address pool for DHCP is an object that contains information about prefix configuration parameters and guidelines that govern the assignment of these prefixes

to requesting routers. If you configured an interface for prefix delegation, the prefix assigned to that interface takes precedence over the prefix or range of prefixes configured at the router level in an IPv6 local pool.

To configure an IPv6 local address pool to be used for DHCPv6 prefix delegation:

1. Enable the IPv6 local address pool for to assign prefixes to the requesting router.

```
host1(config)#ipv6 address-pool local
```

2. Configure the name of the IPv6 local address pool from which the delegating router assigns prefixes to the DHCPv6 client or requesting router.

```
host1(config)#ipv6 local pool dhcpv6pd_pool
```



NOTE: You must enable the IPv6 local address pool feature to be able to configure IPv6 local address pools.

3. Specify the IPv6 prefix range from which prefixes can be delegated to the DHCPv6 client. You can specify the prefix range in one of the following ways:

- Configure the prefix range by specifying an IPv6 prefix and the length of the prefix to be delegated. This prefix length is also called the assigned prefix length.

```
host1(config-v6-local)#prefix 2002:2002::/32 48
```

In this case, the starting and ending prefixes of the range are implicitly specified. In this example, the start of the range is 2002:2002::/48 and the end of the range is 2002:2002:ffff::/48. All prefixes assigned from this range have 48 as the prefix length.

- Alternatively, configure the prefix range by specifying the starting and ending IPv6 prefixes of the range.

```
host1(config-v6-local)#prefix 3003:3003::/56 3003:3003:0:1000::/56
```

In this case, the starting and ending prefixes of the range are explicitly specified. In the preceding example, a prefix range is configured with 16 prefixes that can be allocated to clients. All prefixes assigned from this range have 56 as the prefix length. When you specify the prefix range in this way, you must ensure that the starting and ending prefixes are of the same length.

4. Specify the time period when the requesting router can use the prefix. You can configure a preferred lifetime or a valid lifetime for the requesting router to use when you configure the prefix range. If no lifetime is specified when you configure the prefix range, the default lifetime of 1 day is assigned.



NOTE: The preferred lifetime must be less than or equal to the valid lifetime.

- Specify the number of days and, optionally, the number of hours, minutes, and seconds. You cannot specify a lifetime of zero (that is, you cannot set the days, hours, minutes, and seconds fields all to zero).

```
host1(config-v6-local)#prefix 5005:5005::/32 48 preferred 1 2 3 4
```

In this example, the preferred lifetime is set to 1 day, 2 hours, 3 minutes, and 4 seconds. Because the valid lifetime is not configured, the default value of 1 day is assigned.

- Use the **infinite** keyword to specify a lifetime that does not expire.

```
host1(config-v6-local)#prefix 5005:5005::/32 48 valid infinite
```

In this example, the period for which the prefix remains valid indefinitely for the requesting router to use after it has been delegated by the DHCPv6 server. In this case, the preferred lifetime is set to 1 day by default.

5. Specify the IPv6 address of the DNS servers to be returned to the client. You can configure a primary and secondary DNS server. The DNS server addresses are returned to the client in DHCPv6 responses as part of the DNS Recursive Name Server option.

```
host1(config-v6-local)#dns-server 3001::1 3001::2
```

If the DNS server is not configured in the IPv6 local address pool, the DNS server configured on the DHCPv6 local server is used to delegate prefixes. However, if DNS servers are configured both in the IPv6 local pool and on the DHCPv6 local server, the values configured in the IPv6 local pool take precedence.

6. Specify the name of a DNS domain in the IPv6 local pool to be returned to clients in the DHCPv6 responses as part of the Domain Search List option. The client uses this domain name for DNS resolution. You can specify a maximum of four DNS domains for an IPv6 local pool's search list.

```
host1(config-v6-local)#dns-domain-search test1.com
host1(config-v6-local)#dns-domain-search test2.com
```

You can configure one domain name per line. Enter the command on separate lines to configure additional domain names.

7. Set certain prefixes to be excluded from being allocated to the requesting router. You can exclude those addresses that are assigned to local interfaces. You can exclude specific prefixes or a range of prefixes from delegation to clients.

```
host1(config-v6-local)#exclude-prefix 5005:5005:2::/48 5005:5005:a::/48
```

In this example, all prefixes between the starting prefix of the range, 5005:5005:2::/48, and the ending prefix of the range, 5005:5005:a::/48 are excluded from allocation to clients.

8. Map the domain name to the IPv6 local address pool, which is used for prefix delegation. If the authentication server returns the prefix pool name in the Framed-Ipv6-Pool attribute of the RADIUS-Accept-Request message, this value overrides the IPv6 local pool configured using the **ipv6-prefix-pool-name** command.

```
host1(config)#aaa domain-map westford.com
host1(config-domain-map)#ipv6-prefix-pool-name local_addr_pool
```

For more information about mapping domain names to the IPv6 local address pool, see *ipv6-prefix-pool-name*.

- Related Documentation**
- *aaa domain-map*
 - *dns-domain-search*
 - *dns-server*
 - *exclude-prefix*
 - *prefix*
 - *ipv6 address-pool local*
 - *ipv6 local pool*
 - *ipv6-prefix-pool-name*

Example: Limiting the Number of Prefixes Used by DHCPv6 Clients

If you configure a very large prefix range in an IPv6 local address pool, the number of prefixes that can be used from that range by DHCPv6 clients is limited to 1048576.

Consider the following example in which an IPv6 local address pool, `largePrefixRange`, is configured. The prefix range is specified by the starting prefix and its length as `3003:3003::/32`.

```
host1(config)#ipv6 local pool largePrefixRange
host1(config-v6-local)#prefix 3003:3003::/32 64
host1(config-v6-local)#end
```

The **Total** field of the output of the following `show ipv6 local pool largePrefixRange` and `show ipv6 local pool` commands indicates the number of prefixes that can be allocated to DHCPv6 clients: 1048576.

```
host1#show ipv6 local pool largePrefixRange
```

```
Pool : largePrefixRange
-----
Utilization : 0
      Start                               End                               Total    In Use
-----
3003:3003::/64          3003:3003:ffff:ffff::/64          1048576    0
      Start                               Exclude  Util    Preferred    Valid
      Start                               Exclude  Util    Lifetime    Lifetime
-----
3003:3003::/64          0           0       1 day       1 day
```

```
host1#show ipv6 local pool
```

```
IPv6 Local Address Pools
-----
Pool                               Start                               End
-----
largePrefixRange 3003:3003::/64          3003:3003:ffff:ffff::/64
Pool                               Total    In Use
-----
largePrefixRange 1048576    0
```

Related Documentation • *show ipv6 local pool*

Example: Using DHCPv6 Local Address Pools for Prefix Delegation over non-PPP Links

When a customer premises equipment (CPE) or requesting router and the provider edge (PE) router are connected using a PPP link, one of the following pool names is used to determine the IPv6 local address pool to be used for DHCPv6 Prefix Delegation to the CPE:

- The pool name returned by the RADIUS server in the Framed-IPv6-Pool attribute
- The pool name configured in the AAA domain map

However, for a CPE that is connected to the PE router using a non-PPP link, such as Ethernet, VLAN, or S-VLAN, the method for authentication of clients for DHCPv6 Prefix Delegation is not available in JunosE Release 10.1.x. In such cases, you can select the pool to be used for delegation of prefixes to the CPE by ensuring that the address of the interface over which the DHCPv6 request is received corresponds to any one of the prefix ranges in the configured local address pool.

The following example shows how you can configure an interface with an IPv6 address that matches a prefix configured in an IPv6 local address pool to enable allocation of prefixes from the configured pool for client requests over non-PPP links.

```
! Configure an IPv6 local address pool named example. Specify the IPv6 prefix
! range from which prefixes can be delegated to DHCPv6 clients by specifying an
! IPv6 prefix and the assigned prefix length. Configure the prefix 4004:4004::/48
! to be excluded from being allocated to the requesting client. Exit the IPv6 Local
! Pool Configuration mode.
host1(config)#ipv6 local pool example
host1(config-v6-local)#prefix 4004:4004::/32 48
host1(config-v6-local)#exclude-prefix 4004:4004::/48
host1(config-v6-local)#exit
!
! Create a loopback interface with the IPv6 address matching that of a prefix range
! configured in the example local pool. Exit the Interface Configuration mode.
host1(config)#interface loopback 1
host1(config-if)#ipv6 address 4004:4004::1/48
host1(config-if)#exit
!
! Create a Gigabit Ethernet interface and assign VLAN as the encapsulation
! method. Exit the Interface Configuration mode.
host1(config)#interface gigabitEthernet 2/1/4
host1(config-if)#encapsulation vlan
host1(config-if)#exit
!
! Create a VLAN subinterface, assign a loopback address to it, and enable
! IPv6 Neighbor Discovery. Exit the Interface Configuration mode.
host1(config)#interface gigabitEthernet 2/1/4.100
host1(config-if)#vlan id 100
host1(config-if)#ipv6 unnumbered loopback 1
host1(config-if)#ipv6 nd
host1(config-if)#exit
```


When the PE router receives a request for DHCPv6 Prefix Delegation over the gigabit Ethernet interface 2/1/4.100, prefixes are allocated to the client from the example local pool. In this example, the local pool to use for allocation of prefixes is selected based on the IPv6 address of the interface over which the request is received.

- Related Documentation**
- *dns-domain-search*
 - *dns-server*
 - *exclude-prefix*
 - *interface*
 - *interface loopback*
 - *ipv6 address*
 - *ipv6 nd*
 - *ipv6 unnumbered*
 - *prefix*
 - *ipv6 address-pool local*
 - *ipv6 local pool*
 - *ipv6-prefix-pool-name*
 - *vlan id*

Configuring IPv6 Neighbor Discovery Local Address Pools

The IPv6 local address pool for Neighbor Discovery router advertisements is an object that contains information about prefix configuration parameters and guidelines that govern the assignment of these prefixes to requesting PPPv6 subscribers. If you configured an interface for the Neighbor Discovery router advertisements prefix, the prefix assigned to that interface takes precedence over the prefix or range of prefixes configured at the router level in an IPv6 local address pool.

To configure an IPv6 local address pool to be used for Neighbor Discovery router advertisements:

1. Enable the IPv6 local address pool for Neighbor Discovery router advertisements to assign prefixes to the requesting PPPv6 subscribers.

```
host1(config)#ipv6 address-pool ndra
```

2. Configure the name of the IPv6 local address pool for Neighbor Discovery router advertisements from which the delegating router assigns prefixes to the Neighbor Discovery router advertisements client or requesting router.

```
host1(config)#ipv6 local ndra-pool ndra-pool1
```



NOTE: You must enable the IPv6 local address for Neighbor Discovery router advertisements feature to be able to configure IPv6 local address pools.

3. Specify the IPv6 Neighbor Discovery router advertisements prefix range from which prefixes can be allocated to the Neighbor Discovery router advertisements client.

Configure the prefix range by specifying the starting and ending IPv6 prefixes of the range. The prefix length should be /64. Any attempt to configure a prefix length other than /64 will show an error message.

```
host1(config-v6-NdRa)#ndraprefix 3003:3003::/64 3003:3003:0:1000::/64
```

4. Set certain prefixes for Neighbor Discovery router advertisements to be excluded from being allocated to the requesting PPPv6 subscribers. You can exclude addresses that are assigned to local interfaces. You can exclude specific prefixes or a range of prefixes from allocation to clients.

```
host1(config-v6-NdRa)#exclude-ndraprefix 5005:5005:2::/64 5005:5005:a::/64
```

In this example, all prefixes between the starting prefix of the range 5005:5005:2::/64, and the ending prefix of the range 5005:5005:a::/64, are excluded from allocation to clients.

5. Map the domain name to the IPv6 local address pool, which is used for Neighbor Discovery router advertisements. If the authentication server returns the prefix pool name in the Framed-Ipv6-Pool attribute of the RADIUS-Accept-Request message, this value overrides the IPv6 local pool configured using the **ipv6-ndra-pool-name** command.

```
host1(config)#aaa domain-map westford.com
host1(config-domain-map)#ipv6-ndra-pool-name local_addr_pool
```

For more information about mapping domain names to the IPv6 local address pool, see *ipv6-ndra-pool-name*.

Related Documentation

- [IPv6 Prefix Allocation Using Neighbor Discovery Router Advertisements from IPv6 Address Pools Overview on page 57](#)
- *aaa dhcpv6-ndra-pool override*
- *exclude-ndraprefix*
- *ipv6 address-pool ndra*
- *ipv6 local ndra-pool*
- *ndraprefix*

Overriding AAA to Perform IPCP and IPv6CP Negotiations Based on RADIUS-Returned Attributes

Normally, AAA performs IPCP and IPv6CP negotiations based on the RADIUS attributes returned from the RADIUS server and the subscriber domain map configurations. You can override the standard IPCP and IPv6CP negotiations by using the **aaa radius-override-ncp-negotiation** command, which allows you to enable or disable IPCP and IPv6CP negotiations only when one of the following RADIUS attributes is returned in the Access-Accept message:

- Framed-Ip-Address [8]
- Framed-Pool [88]
- Framed-Ipv6-Prefix [97]
- Ipv6-NdRa-Prefix [26-129]
- Framed-Ipv6-Pool [100]
- Ipv6-Ndra-Pool [26-157]
- Framed-Interface-Id [96]
- Delegated-Ipv6-Prefix [123]

To enable IPCP and IPv6CP negotiations for IP and IPv6 clients based on RADIUS-returned attributes:

- Issue the **aaa radius-override-ncp-negotiation** command with the **enable** keyword in Global Configuration mode.

```
host1(config)#aaa radius-override-ncp-negotiation enable
```

To disable IPCP and IPv6CP negotiations for IP and IPv6 clients based on RADIUS-returned attributes:

- Issue the **aaa radius-override-ncp-negotiation** command with the **disable** keyword in Global Configuration mode.

```
host1(config)#aaa radius-override-ncp-negotiation disable
```

Use the **no** version to disable IPCP and IPv6CP negotiations based on RADIUS-returned attributes, which is the default behavior. IPCP and IPv6CP negotiations are initiated on the basis of RADIUS-returned attributes and configurations done on the subscriber domain map.

Related Documentation

- [Understanding IPCP and IPv6CP Negotiations for IPv4 and IPv6 Clients Based on RADIUS-Returned Attributes on page 60](#)
- [Monitoring the Status of the Override Feature to Initiate IPCP and IPv6CP Negotiations Based on RADIUS-Returned Attributes on page 175](#)
- *aaa radius-override-ncp-negotiation*

Sending the Acct-Start and Acct-Update Messages After NCP Negotiation Is Completed

You can use the **aaa accounting delay-start** command to delay the sending of the Acct-Start message from the Broadband Remote Access Server (B-RAS) to the RADIUS server until IPCP or IPv6CP negotiation is completed. When Network Control Protocol (NCP) negotiation is completed, IP notifies authentication, authorization, and accounting (AAA) to set a user IP address or user interface ID (in case of Point-to-Point Protocol version 6 [PPPoE] subscribers). Only then does the B-RAS send the Acct-Start message to the RADIUS server. This ensures that the Framed-Interface-Id [96] or Framed-IP-Address [8] RADIUS attributes that are negotiated during Internet Protocol Control Protocol (IPCP) or Internet Protocol version 6 Control Protocol (IPv6CP) are present in the Acct-Start message, thereby reducing the number of accounting messages from two to one for each subscriber.

To delay the sending of the Acct-Start message from the B-RAS to the RADIUS server:

- Issue the **aaa accounting delay-start** command in Global Configuration mode.

```
host1(config)#aaa accounting delay-start enable
```

Use the **no** version to immediately send the Acct-Start message from the B-RAS to the RADIUS server when AAA receives authentication-grant permission from the RADIUS server.

You can use the **aaa accounting immediate-update-framed-ipv6-interfaceId-negotiation** command to send the Acct-Update message with the Framed-Interface-Id [96] RADIUS attributes from the B-RAS to the RADIUS server for PPPoE users immediately after NCP negotiation is completed.

To send the Acct-Update message with the Framed-Interface-Id [96] RADIUS attributes from the B-RAS to the RADIUS server:

- Issue the **aaa accounting immediate-update-framed-ipv6-interfaceId-negotiation** command in Global Configuration mode.

```
host1(config)#aaa accounting immediate-update-framed-ipv6-interfaceId-negotiation enable
```

Use the **no** version to immediately send the Acct-Update message without the Framed-Interface-Id [96] RADIUS attributes from the B-RAS to the RADIUS server after IP notifies AAA to set a user IP address or user interface ID.

Related Documentation

- [Sending the Acct-Start and Acct-Update Messages After NCP Negotiation Is Completed Overview on page 63](#)
- [Monitoring AAA Accounting Configuration on page 115](#)
- [Monitoring AAA Configuration to Send the Acct-Start Message After NCP Negotiation Is Completed on page 117](#)
- [Monitoring AAA Configuration to Send the Acct-Update Message After NCP Negotiation Is Completed on page 117](#)

- *aaa accounting delay-start*
- *aaa accounting immediate-update-framed-ipv6-interfaceid-negotiation*

CHAPTER 3

Monitoring and Troubleshooting Remote Access

Use the commands in this chapter to set baselines for and to monitor remote access.

- [Setting Baselines for Remote Access on page 113](#)
- [How to Monitor PPP Interfaces on page 115](#)
- [Monitoring AAA Accounting Configuration on page 115](#)
- [Monitoring AAA Configuration to Send the Acct-Start Message After NCP Negotiation Is Completed on page 117](#)
- [Monitoring AAA Configuration to Send the Acct-Update Message After NCP Negotiation Is Completed on page 117](#)
- [Monitoring AAA Accounting Default on page 118](#)
- [Monitoring the AAA Accounting Interval on page 118](#)
- [Monitoring AAA-Specific Virtual Router Groups on page 118](#)
- [Monitoring the Default AAA Authentication Method List on page 119](#)
- [Monitoring AAA Domain Name Stripping for a Domain Per Virtual Router on page 119](#)
- [Monitoring Mapping Between User Domains and Virtual Routers on page 120](#)
- [Monitoring Tunnel Subscriber Authentication on page 122](#)
- [Monitoring Routing Table Address Lookup on page 123](#)
- [Monitoring the AAA Model on page 123](#)
- [Monitoring IP Addresses of Primary and Secondary DNS and WINS Name Servers on page 123](#)
- [Monitoring AAA Profile Configuration on page 124](#)
- [Monitoring Statistics about the RADIUS Route-Download Server on page 125](#)
- [Monitoring Routes Downloaded by the RADIUS Route-Download Server on page 127](#)
- [Monitoring Chassis-Wide Routes Downloaded by the RADIUS Route-Download Server on page 129](#)
- [Monitoring AAA Statistics on page 131](#)
- [Monitoring the Number of Active Subscribers Per Port on page 133](#)
- [Monitoring the Maximum Number of Active Subscribers Per-Virtual Router on page 133](#)

- [Monitoring Session Timeouts on page 134](#)
- [Monitoring Interim Accounting for Users on the Virtual Router on page 134](#)
- [Monitoring Virtual Router Groups Configured for AAA Broadcast Accounting on page 134](#)
- [Monitoring Configuration Information for AAA Local Authentication on page 135](#)
- [Monitoring AAA Server Attributes on page 136](#)
- [Monitoring the COPS Layer Over SRC Connection on page 138](#)
- [Monitoring Statistics About the COPS Layer on page 140](#)
- [Monitoring Local Address Pool Aliases on page 142](#)
- [Monitoring Local Address Pools on page 143](#)
- [Monitoring Local Address Pool Statistics on page 144](#)
- [Monitoring Shared Local Address Pools on page 145](#)
- [Monitoring the Routing Table on page 145](#)
- [Monitoring the B-RAS License on page 146](#)
- [Monitoring the RADIUS Server Algorithm on page 146](#)
- [Monitoring RADIUS Override Settings on page 146](#)
- [Monitoring the RADIUS Rollover Configuration on page 147](#)
- [Monitoring RADIUS Server Information on page 147](#)
- [Monitoring RADIUS Services Statistics on page 149](#)
- [Monitoring RADIUS SNMP Traps on page 153](#)
- [Monitoring RADIUS Accounting for L2TP Tunnels on page 153](#)
- [Monitoring RADIUS UDP Checksums on page 153](#)
- [Monitoring RADIUS Server IP Addresses on page 154](#)
- [Monitoring the RADIUS Attribute Used for IPv6 Neighbor Discovery Router Advertisements on page 154](#)
- [Monitoring the RADIUS Attribute Used for DHCPv6 Prefix Delegation on page 154](#)
- [Monitoring Duplicate IPv6 Prefixes on page 154](#)
- [Monitoring Duplicate IPv6 Prefixes in the AAA User Profile Database on page 155](#)
- [Monitoring the SRC Client Connection Status on page 155](#)
- [Monitoring SRC Client Connection Statistics on page 158](#)
- [Monitoring the SRC Client Version Number on page 159](#)
- [Monitoring the SRC Option on page 160](#)
- [Monitoring Subscriber Information on page 161](#)
- [Monitoring Application Terminate Reason Mappings on page 167](#)
- [Monitoring IPv6 Local Pools for DHCP Prefix Delegation By All Configured Pools on page 169](#)
- [Monitoring IPv6 Local Pools for DHCP Prefix Delegation By Pool Name on page 170](#)
- [Monitoring IPv6 Local Pool Statistics for DHCP Prefix Delegation on page 171](#)

- [Monitoring IPv6 Local Pools for Neighbor Discovery Router Advertisements for all Configured Pools on page 172](#)
- [Monitoring IPv6 Local Pools for Neighbor Discovery Router Advertisements by Pool Name on page 173](#)
- [Monitoring IPv6 Local Pool Statistics for Neighbor Discovery Router Advertisements Allocation of Prefixes on page 174](#)
- [Monitoring the Status of the Override Feature to Initiate IPCP and IPv6CP Negotiations Based on RADIUS-Returned Attributes on page 175](#)

Setting Baselines for Remote Access

You can set baseline statistics using the **baseline** commands. The router implements the baseline by reading and storing the statistics at the time the baseline is set and then subtracting this baseline when you retrieve baseline-relative statistics.

Issue the **delta** keyword with the **show aaa statistics** command to show baselined statistics.

1. [Setting a Baseline for AAA Statistics on page 113](#)
2. [Setting a Baseline for AAA Route Downloads on page 113](#)
3. [Setting a Baseline for COPS Statistics on page 114](#)
4. [Setting a Baseline for Local Address Pool Statistics on page 114](#)
5. [Setting a Baseline for RADIUS Statistics on page 114](#)
6. [Setting the Baseline for SRC Statistics on page 114](#)

Setting a Baseline for AAA Statistics

Purpose Set a baseline for all AAA statistics.

Action Issue the **baseline aaa** command:

```
host1#baseline aaa
```

There is no **no** version.

Setting a Baseline for AAA Route Downloads

Purpose Set a baseline for route downloads.

Action • Issue the **baseline aaa route-download** command for IPv4 routes:

```
host1#baseline aaa route-download
```

• Issue the **baseline aaa route-download ipv6** command for IPv6 routes:

```
host1#baseline aaa route-download ipv6
```

There is no **no** version.

Setting a Baseline for COPS Statistics

Purpose Set a baseline for COPS statistics.

Action Issue the **show cops statistics** command:

```
host1#show cops statistics
```

There is no **no** version.

Setting a Baseline for Local Address Pool Statistics

Purpose Set a baseline for local address pool statistics.

Action Issue the **show local pool statistics** command:

```
host1#show local pool statistics
```

There is no **no** version.

Setting a Baseline for RADIUS Statistics

Purpose Set a baseline for RADIUS statistics.

Action Issue the **show radius statistics** command:

```
host1#show radius statistics
```

There is no **no** version.

Setting the Baseline for SRC Statistics

Purpose Set a baseline for SRC statistics.

Action Issue the **show ssc statistics** command:

```
host1#show ssc statistics
```

There is no **no** version.

- Related Documentation**
- *baseline aaa*
 - *baseline aaa route-download*
 - *baseline cops*
 - *baseline local pool*
 - *baseline radius*
 - *baseline ssc*

How to Monitor PPP Interfaces

Purpose Monitor PPP interfaces.

Action Use the following commands:

- **show ppp interface summary**
- **show ppp interface** <selective control>

For details on the **show ppp** commands, see *JunosE Link Layer Configuration Guide*.

You can use the output filtering feature of the **show** command to include or exclude lines of output based on a text string you specify. For details, see *JunosE System Basics Configuration Guide*.



NOTE: AAA and RADIUS statistics are not preserved across a warm restart when stateful SRP Switchover is enabled.

- Related Documentation**
- *Monitoring PPP Interfaces*
 - *Monitoring Multilinked and Nonmultilinked PPP Interfaces*

Monitoring AAA Accounting Configuration

Purpose Display the AAA accounting configuration.

Action To display the **show aaa accounting** command:

```
host1:vrXyz7#show aaa accounting
```

```
Accounting duplication set to router vrXyz25
Broadcast accounting uses group groupXyzCompany20
send acct-stop on AAA access deny is enabled
send acct-stop on authentication server access deny is disabled
user-acct-interval 0
service-acct-interval 0
send immediate-update is enabled
send interim-update is enabled
send delay-start is disabled
send immediate-update-ipv6 is disabled
statistics volume-time
```

Meaning Table 12 on page 115 lists the **show aaa accounting** command output fields.

Table 12: show aaa accounting Output Fields

Field Name	Field Description
Accounting duplication	Name of the virtual router whose accounting server receives duplicate accounting records

Table 12: show aaa accounting Output Fields (*continued*)

Field Name	Field Description
Broadcast accounting	Name of the virtual router group whose accounting servers receive broadcast accounting records
send acct-stop on AAA access deny	Specifies whether the Acct-Stop message is enabled or disabled when access to the AAA server is denied
send acct-stop on authentication server access deny	Specifies whether the Acct-Stop message is enabled or disabled when access to the authentication server is denied
user-acct-interval	Number of minutes between interim accounting updates for users on the virtual router
service-acct-interval	Number of minutes between interim accounting updates for services created by the Service Manager feature
send immediate-update	Specifies whether the sending of the Acct-Update message on receipt of the response to the Acct-Start message is enabled or disabled
send interim-update	Specifies whether the periodic interim update at the configured user accounting interval is enabled or disabled
send delay-start	Specifies whether the delay in the sending of the Acct-Start message until NCP negotiation is completed is enabled or disabled
send immediate-update-ipv6	Specifies whether the sending of the Acct-Update message with the Framed-Interface-Id [96] attribute after NCP negotiation is completed is enabled or disabled
statistics	Specifies whether a full set of statistics or only an uptime status is collected for the sessions managed by AAA. This value can be volume-time or time.

Related Documentation

- [Configuring Interim, Broadcast, and Policy-Based Accounting in Virtual Router Groups and PPP Profiles on page 89](#)
- [Sending the Acct-Start and Acct-Update Messages After NCP Negotiation Is Completed on page 108](#)
- *show aaa accounting*

Monitoring AAA Configuration to Send the Acct-Start Message After NCP Negotiation Is Completed

Purpose	Display authentication, authorization, and accounting (AAA) configuration to delay the sending of the Acct-Start message until IPCP or IPv6CP negotiation is completed.
Action	To display the status of delaying the sending of the Acct-Start message: <pre>host1#show aaa send-delay-start</pre> <p>disabled</p>
Related Documentation	<ul style="list-style-type: none"> • Sending the Acct-Start and Acct-Update Messages After NCP Negotiation Is Completed Overview on page 63 • Sending the Acct-Start and Acct-Update Messages After NCP Negotiation Is Completed on page 108 • <i>aaa accounting delay-start</i> • <i>show aaa accounting</i> • <i>show aaa send-delay-start</i>

Monitoring AAA Configuration to Send the Acct-Update Message After NCP Negotiation Is Completed

Purpose	Display authentication, authorization, and accounting (AAA) configuration to send the Acct-Update message with the Framed-Interface-Id [96] RADIUS attribute after NCP negotiation is completed. In a single-stack environment, the Acct-Update message is sent after IPv6CP negotiation is completed. In a dual-stack environment, the Acct-Update message is sent after both IPCP and IPv6CP negotiations are completed.
Action	To display the status of the sending of the Acct-Update message after NCP negotiation is completed: <pre>host1#show aaa send-immediate-update-ipv6</pre> <p>disabled</p>
Related Documentation	<ul style="list-style-type: none"> • Sending the Acct-Start and Acct-Update Messages After NCP Negotiation Is Completed Overview on page 63 • Sending the Acct-Start and Acct-Update Messages After NCP Negotiation Is Completed on page 108 • <i>aaa accounting immediate-update-framed-ipv6-interfaceId-negotiation</i> • <i>show aaa accounting</i> • <i>show aaa send-immediate-update-ipv6</i>

Monitoring AAA Accounting Default

Purpose Display the AAA accounting default method for a subscriber type.

You can view the method used for ATM 1483, IPsec, PPP, RADIUS relay server, and tunnel subscribers, and IP subscriber management interfaces.

Action To display the default AAA accounting method:

```
host1#show aaa accounting tunnel default
radius
```

Related Documentation

- *show aaa accounting default*

Monitoring the AAA Accounting Interval

Purpose Display the accounting interval.

Action To display the accounting interval:

```
host1#show aaa accounting interval
acct-interval (for PPP Clients) 10
```

Related Documentation

- *show aaa accounting interval*

Monitoring AAA-Specific Virtual Router Groups

Purpose Display the names of a specific virtual router group or of all virtual router groups configured on the router, and of the virtual routers making up the groups.

Action To display the names of a specific virtual router group or of all virtual router groups configured on the router. Display the virtual routers making up the groups:

```
host1#show aaa accounting vr-group
vr-group groupXyzCompany10:
aaa virtual-router 1 vrXyzA
aaa virtual-router 1 vrXyzA interim-accounting enabled
aaa virtual-router 1 vrXyzA policy-accounting disabled
aaa virtual-router 2 vrXyzB
aaa virtual-router 2 vrXyzB interim-accounting enabled
aaa virtual-router 2 vrXyzB policy-accounting disabled
aaa virtual-router 3 vrXyzC
aaa virtual-router 3 vrXyzC interim-accounting enabled
aaa virtual-router 3 vrXyzC policy-accounting disabled
aaa virtual-router 4 vrXyzD
aaa virtual-router 4 vrXyzD interim-accounting enabled
aaa virtual-router 4 vrXyzD policy-accounting disabled
vr-group groupXyzCompany20:
aaa virtual-router 1 vrXyzP
aaa virtual-router 1 vrXyzP interim-accounting enabled
aaa virtual-router 1 vrXyzP policy-accounting disabled
aaa virtual-router 2 vrXyzQ
aaa virtual-router 2 vrXyzQ interim-accounting enabled
```

```

aaa virtual-router 2 vrXyzQ policy-accounting disabled
aaa virtual-router 3 vrXyzR
aaa virtual-router 3 vrXyzR interim-accounting enabled
aaa virtual-router 3 vrXyzR policy-accounting disabled
aaa virtual-router 4 vrXyzS
aaa virtual-router 4 vrXyzS interim-accounting enabled
aaa virtual-router 4 vrXyzS policy-accounting disabled

```

Meaning [Table 13 on page 119](#) lists the **show aaa accounting vr-group** command output fields.

Table 13: show aaa accounting vr-group Output Fields

Field Name	Field Description
vr-group	Name of the virtual router group

- Related Documentation**
- [Configuring AAA Broadcast Accounting on page 68](#)
 - [Configuring Interim, Broadcast, and Policy-Based Accounting in Virtual Router Groups and PPP Profiles on page 89](#)
 - *show aaa accounting vr-group*

Monitoring the Default AAA Authentication Method List

Purpose Display the default AAA authentication method list for a subscriber type. You can view the method list used for ATM 1483 subscribers, IPsec subscribers, IP subscriber management interfaces, PPP subscribers, RADIUS relay subscribers, and tunnel subscribers.

For example, you can verify that the local authentication method is configured for PPP subscribers.

Action To display the default AAA authentication method list for a subscriber type:

```

host1#show aaa authentication ppp default
local none

```

- Related Documentation**
- *show aaa authentication default*

Monitoring AAA Domain Name Stripping for a Domain Per Virtual Router

Purpose Display information about the aaa domain-name stripping functionality per virtual router.

Action To display information about the aaa domain-name stripping functionality per virtual router:

```

host1:vr1(config)#show aaa strip-domain
strip-domain is disable
strip-domain domainName delimiter is "@"
strip-domain domainName parse direction is right-to-left

```

Meaning Table 14 on page 120 lists the **show aaa strip-domain** command output fields.

Table 14: show aaa strip-domain Output Fields

Field Name	Field Description
delimiter	Delimiter value configured for the subscriber's domain
domainName	The domain name characteristics configured for the broadband remote access subscriber per virtual router
disable	The domain name stripping functionality is disabled for the virtual router
enable	The domain name stripping functionality is enabled for the virtual router
left-to-right	The parsing direction configured for stripping the domain name at the virtual router is left-to-right
right-to-left	The parsing direction configured for stripping the domain name at the virtual router is right-to-left

- Related Documentation**
- *aaa domain-map*
 - *ppp authentication*
 - *show aaa delimiters*
 - *show aaa strip-domain*

Monitoring Mapping Between User Domains and Virtual Routers

Purpose Display the mapping between user domains and virtual routers.

The following keywords have significance when used as user domains:

- **none**—All client requests with no user domain name are associated with the virtual router mapped to the **none** entry
- **default**—All client requests with a domain present that have no map are associated with the virtual router mapped to the **default** entry

Action To display the mapping between user domains and virtual routers:

```
host1#show aaa domain-map
Domain: lac-tunnel; auth-router-name: lac;
ip-router-name: default; ipv6-router-name: default
Tunnel      Tunnel  Tunnel  Tunnel  Tunnel
Tag         Tunnel Peer  Source  Type    Medium  Password  Tunnel Id
-----
5           192.168.1.1 <null>  12tp    ipv4     welcome  lac-tunnel
```


Tunnel Tag	Tunnel Client Name	Tunnel Server Name	Tunnel Preference	Tunnel Max Sessions	Tunnel RWS
5	lac	boston	5	0	4
Tunnel Tag	Tunnel Virtual Router	Tunnel Failover Resync	Tunnel Switch Profile	Tunnel Tx Speed Method	
5	<null>	silent failover	denver	qos	

Meaning Table 15 on page 121 lists the **show aaa domain-map** command output fields.

Table 15: show aaa domain-map Output Fields

Field Name	Field Description
Domain	Name of the domain
auth-router-name	Access virtual router to which user domain name is mapped
ip-router-name	IPv4 virtual router to which user domain name is mapped
router-mask	IP mask of the local interface
tunnel-group	Name of the tunnel group assigned to the domain map
ipv6-router-name	IPv6 virtual router to which user domain name is mapped
local-interface	Interface information to use on the local (E Series) side of the subscriber's interface
ipv6-local-interface	IPv6 interface information to use on the local (E Series) side of the subscriber's interface
poolname	Local address pool from which the router allocates addresses for this domain
IP hint	IP hint is enabled
strip-domain	Strip domain is enabled
override-username	Single username used for all users from a domain in place of the values received from the remote client
override-password	Single password used for all users from a domain in place of the values received from the remote client
Tunnel Tag	Tag that identifies the tunnel

Table 15: show aaa domain-map Output Fields (*continued*)

Field Name	Field Description
Tunnel Peer	Destination address of the tunnel
Tunnel Source	Source address of the tunnel
Tunnel Type	L2TP
Tunnel Medium	Type of medium for the tunnel; only IPv4 is supported
Tunnel Password	Password for the tunnel
Tunnel Id	ID of the tunnel
Tunnel Client Name	Hostname that the LAC sends to the LNS when communicating to the LNS about the tunnel
Tunnel Server Name	Hostname expected from the peer (the LNS) when during tunnel startup
Tunnel Preference	Preference level for the tunnel
Tunnel Max Sessions	Maximum number of sessions allowed on a tunnel
Tunnel RWS	L2TP receive window size (RWS) for a tunnel on the LAC; displays either the configured value or the default behavior, which is indicated by system chooses (that is, the RWS value is set as 64)
Tunnel Virtual Router	Name of the virtual router to map to the user domain name
Tunnel Failover Resync	L2TP peer resynchronization method
Tunnel Switch Profile	Name of the L2TP tunnel switch profile
Tunnel Tx Speed Method	Method that the router uses to calculate the transmit connect speed of the subscriber's access interface: static layer2, dynamic layer2, qos, actual, not set

Related Documentation

- *show aaa domain-map*

Monitoring Tunnel Subscriber Authentication

Purpose Verify configuration of tunnel subscriber authentication. When authentication is enabled, the output indicates this configuration. When authentication is disabled, the output presents no information about the configuration.

Action To display tunnel subscriber authentication configuration:

```
host1#show aaa domain-map
Domain: tunnel.com; auth-router-name: default; ip-router-name: default
ipv6-router-name: default; tunnel-subscriber authentication: enable
```

Meaning Authentication is enabled.

Related Documentation

- *show aaa domain-map*

Monitoring Routing Table Address Lookup

Purpose Display whether the routing table address lookup or duplicate address check is enabled or disabled.

Action To display whether the routing table address lookup or duplicate address check is enabled or disabled:

```
host1#show aaa duplicate-address-check
enabled
```

Related Documentation

- *show aaa duplicate-address-check*

Monitoring the AAA Model

Purpose Display the AAA model.

Action To display the AAA model:

```
host1#show aaa model
aaa model: old model
```

Related Documentation

- *show aaa model*

Monitoring IP Addresses of Primary and Secondary DNS and WINS Name Servers

Purpose Display the IP addresses of the primary and secondary DNS and WINS name servers.

Action To display the IP addresses of the primary and secondary DNS and WINS name servers:

```
host1#show aaa name-servers
Name Server Addresses (for PPP Clients):
primary DNS Addr          10.2.3.4
secondary DNS Addr        10.6.7.8
primary NBNS (WINS) Addr  10.22.33.44
secondary NBNS (WINS) Addr 10.66.77.88
```

Meaning The IP addresses of DNS and WINS name servers are displayed.

Related Documentation • [show aaa name-servers](#)

Monitoring AAA Profile Configuration

Purpose Display the configuration of all AAA profiles or of a specific profile.

Action To display the configuration of all AAA profiles or of a specific profile:

```
host1#show aaa profile name PreAuth1
preAuth1:
  atm nas-port-type: ADLSL-CAP
  ethernet nas-port-type: Cable
  profile-service-description: xyzService
  pre-authenticate
  allow xyz.com
  deny default
  translate xyz1.com abc.com
  aaaPerProfileName:aaaProfile1
  radiusPerProfileName:radiusProfile1
```

Meaning [Table 16 on page 124](#) lists the **show aaa profile** command output fields.

Table 16: show aaa profile Output Fields

Field Name	Field Description
atm nas-port-type	Configuration of NAS-Port-Type attribute for ATM interfaces
ethernet nas-port-type	Configuration of NAS-Port-Type attribute for Ethernet interfaces
profile-service-description	Description configured in the Service-Description attribute
pre-authenticate	Indicates that subscriber preauthentication is configured for the profile
allow	One or more domain names that are allowed access to AAA authentication
deny	One or more domain names that are denied access to AAA authentication
translate	Original domain name and the name to which it is mapped for domain map lookup
aaaPerProfileName	Name of the AAA per-profile
radiusPerProfileName	Name of the RADIUS per-profile

Related Documentation • *show aaa profile*

Monitoring Statistics about the RADIUS Route-Download Server

Purpose Display statistics about the RADIUS route-download server configuration.

- Use the optional **statistics** keyword to display information about the RADIUS route download server operation.
- Use the optional **delta** keyword to show baselined statistics.

Action To display information about the RADIUS route-download server operation for IPv4 routes:

```
host1#show aaa route-download
AAA Route Downloader:    configured in virtual router default
Download Interval:       720 minutes
Retry Interval:          10 minutes
Default Cost:            2
Default Tag:             0
Base User Name:          <HOSTNAME>
Password:                <DEFAULT>
Synchronization:        <NOT SET>

Status:                  idle
Last Download Attempt:   TUE DEC 19 22:46:47 2006
Last Download Success:   TUE DEC 19 22:46:47 2006
Last Regular Download:   complete
Next Download Scheduled: WED DEC 20 10:46:47 2006
Next Regular Download:   WED DEC 20 10:46:47 2006
```

To display statistics about the RADIUS route-download server configuration for IPv4 routes:

```
host1#show aaa route-download statistics

Total Download Attempts: 2
Successful Downloads:    2
Downloaded Fragments:    3756
Downloaded Routes:       192000
IP Updates:              1
Updated Routes:          96000
Cleared Route Intervals: 0
```

To display information about the RADIUS route-download server operation for IPv6 routes:

```
host1#show aaa route-download ipv6
AAA Route Downloader:    configured in virtual router default
Download Interval:       720 minutes
Retry Interval:          10 minutes
Default Cost:            2
Default Tag:             0
Base User Name:          <HOSTNAME>
Password:                <DEFAULT>
Synchronization:        <NOT SET>

Status:                  idle
Last Download Attempt:   TUE DEC 13 2011 00:05:43 UTC
Last Download Success:   TUE DEC 13 2011 00:05:43 UTC
```

```

Last Regular Download:    complete
Next Download Scheduled:  TUE DEC 13 2011 12:05:42 UTC
Next Regular Download:    TUE DEC 13 2011 12:05:42 UTC

```

To display statistics about the RADIUS route-download server configuration for IPv6 routes:

```
host1#show aaa route-download ipv6 statistics
```

```

Total Download Attempts: 3
Successful Downloads:    3
Downloaded Fragments:    30
Downloaded Routes:       240
IP Updates:              2
Updated Routes:          16
Cleared Route Intervals: 0

```

Meaning [Table 17 on page 126](#) lists the **show aaa route-download** command and **show aaa route-download ipv6** command output fields.

Table 17: show aaa route-download Output Fields

Field Name	Field Description
AAA Route Downloader	Virtual router where the RADIUS route-download server is configured
Download Interval	Number of minutes between route downloads
Retry Interval	Number of minutes before retry after a download failure
Default Cost	Default cost of downloaded routes
Default Tag	Default tag for downloaded routes
Base Username	Virtual router used for route-download requests; either <HOSTNAME> or the configured name
Password	Password for route-download requests or <DEFAULT>
Synchronization	Either <NOT SET> or the time that the server starts the route download operation each day
Status	Current status of route-download server; waiting for base router, waiting for IP warmstart, idle, downloading, updating ip, downloading and updating ip, or suspended
Last Download Attempt	Either <NEVER> or the day, date, and time of attempt
Last Download Success	Either <NEVER> or the day, date, and time of success
Last Regular Download	Status of last regular download; either complete or not complete

Table 17: show aaa route-download Output Fields (*continued*)

Field Name	Field Description
Next Download Scheduled	<DOWNLOAD ACTIVE>, <NOT SCHEDULED>, or the day, date, and time of next download
Next Regular Download	Day, date, and time
Total Download Attempts	Number of downloads attempted
Successful Downloads	Number of successful download operations
Downloaded Fragments	Number of downloaded fragments
Downloaded Routes	Number of downloaded routes
IP Updates	Number of IP updates
Updated Routes	Number of updated routes
Cleared Route Intervals	Number of cleared route intervals

Related Documentation

- *show aaa route-download*

Monitoring Routes Downloaded by the RADIUS Route-Download Server

Purpose Display information about the routes that are downloaded by the RADIUS route-download server.

Use the optional **detail** keyword to display more detailed information about the downloaded routes.

Action To display information about the IPv4 static routes that are downloaded by the RADIUS route-download server:

```
host1#show aaa route-download routes
96000 downloaded routes
```

To display detailed information about the IPv4 static routes that are downloaded by the RADIUS route-download server:

```
host1#show aaa route-download routes detail
Prefix/Length    Type      NextHop      Dst/Met  Intf    Tag
-----
192.168.1.1/32   Access-P  255.255.255.255  254/2    nu110   0
192.168.1.5/32   Access-P  255.255.255.255  254/2    nu110   0
192.168.1.9/32   Access-P  255.255.255.255  254/2    nu110   0
192.168.1.13/32  Access-P  255.255.255.255  254/2    nu110   0
192.168.1.17/32  Access-P  255.255.255.255  254/2    nu110   0
192.168.1.21/32  Access-P  255.255.255.255  254/2    nu110   0
```

To display information about the IPv6 routes that are downloaded by the RADIUS route-download server:

```
host1#show aaa route-download ipv6 routes
13 downloaded routes
```

To display detailed information about the IPv6 routes that are downloaded by the RADIUS route-download server:

```
host1#show aaa route-download ipv6 routes detail
Prefix/Length    Type      Dst/Met  Intf     Tag
-----
f001::1/128      Access-P  0/2      null0    0
f002::1/128      Access-P  0/2      null0    0
f002::2/128      Access-P  0/2      null0    0
f002::3/128      Access-P  0/2      null0    0
f002::4/128      Access-P  0/2      null0    0
f003::2/128      Access-P  0/2      null0    0
f004::2/128      Access-P  0/2      null0    0
f005::2/128      Access-P  0/2      null0    0
f006::2/128      Access-P  0/2      null0    0
f007::2/128      Access-P  0/2      null0    0
f008::2/128      Access-P  0/2      null0    0
f009::2/128      Access-P  0/2      null0    0
f00a::2/128      Access-P  0/2      null0    0
```

Meaning [Table 18 on page 128](#) lists the **show aaa route-download routes** command and **show aaa route-download ipv6 routes** command output fields.

Table 18: show aaa route-download routes Output Fields

Field Name	Field Description
downloaded routes	Number of current downloaded routes
Prefix/Length	IP address prefix and mask information for downloaded routes
Type	Type of downloaded routes; Access-P indicates routes downloaded from the RADIUS route-download server
NextHop	IP address of the next hop
Dst/Met	Administrative distance and number of hops for the route
Tag	Tag assigned to downloaded routes
Intf	Interface type and specifier

Related Documentation

- *show aaa route-download routes*
- *show aaa route-download ipv6 routes*

Monitoring Chassis-Wide Routes Downloaded by the RADIUS Route-Download Server

Purpose Display chassis-wide information about routes that are downloaded by RADIUS route-download servers.

Use the optional **detail** keyword to display more detailed information about the downloaded routes.

Use the optional **start** keyword to specify the first router context that you want to display in the output. For example, `aaa:a2` specifies that the display shows a list of router contexts starting with VRF `a2` in virtual router `aaa`.

Action To display chassis-wide information about IPv4 routes that are downloaded by RADIUS route-download servers:

```
host1#show aaa route-download routes global
```

Virtual Router	VRF	Present	Number of Routes
aaa		n	4
aaa	a1	n	4
default		y	4
default	d1	n	4

To display more detailed information about the downloaded IPv4 routes:

```
host1# show aaa route-download routes global detail
```

Virtual Router	VRF	Present	Prefix/Length	Type	NextHop	Dst/Met	Intf	Tag
aaa		n	192.168.1.1/32	Access-P	255.255.255.255	0/2	null0	0
aaa		n	192.168.1.2/32	Access-P	255.255.255.255	0/2	null0	0
aaa		n	192.168.3.1/32	Access-P	255.255.255.255	0/2	null0	0
aaa		n	192.168.4.1/32	Access-P	255.255.255.255	0/2	null0	0
aaa	a1	n	192.168.5.3/32	Access-P	255.255.255.255	0/2	null0	0
aaa	a1	n	192.168.7.1/32	Access-P	255.255.255.255	0/2	null0	0
aaa	a1	n	192.168.7.5/32	Access-P	255.255.255.255	0/2	null0	0
aaa	a1	n	192.168.9.1/32	Access-P	255.255.255.255	0/2	null0	0
default		y	192.168.22.1/32	Access-P	255.255.255.255	0/2	null0	0
default		y	192.168.23.1/32	Access-P	255.255.255.255	0/2	null0	0
default		y	192.168.24.1/32	Access-P	255.255.255.255	0/2	null0	0
default		y	192.168.25.1/32	Access-P	255.255.255.255	0/2	null0	0
default	d1	n	192.168.40.6/32	Access-P	255.255.255.255	0/2	null0	0
default	d1	n	192.168.40.7/32	Access-P	255.255.255.255	0/2	null0	0
default	d1	n	192.168.40.8/32	Access-P	255.255.255.255	0/2	null0	0
default	d1	n	192.168.40.9/32	Access-P	255.255.255.255	0/2	null0	0

To specify the first router context that you want to display in the output:

```
host1#show aaa route-download routes global start aaa:a2
```

Virtual Router	VRF	Present	Number of Routes
default		y	4
default	d1	n	4

To display chassis-wide information about IPv6 routes that are downloaded by RADIUS route-download servers:

```
host1#show aaa route-download ipv6 routes global
```

Virtual Router	VRF	Present	Number of Routes
def		y	3
def	temp	y	1
default		y	13
Context1		n	27
test		n	36

To display more detailed information about the downloaded IPv6 routes:

```
host1# show aaa route-download ipv6 routes global detail
```

Virtual Router	VRF	Present	Prefix/Length	Type	Dst/Met	Intf	Tag
def		y	f00b::2/128	Access-P	0/2	nu110	0
def		y	f00b::3/128	Access-P	0/2	nu110	0
def		y	f00b::4/128	Access-P	0/2	nu110	0
def	temp	y	f00b::1/128	Access-P	0/2	nu110	0
default		y	f001::1/128	Access-P	0/2	nu110	0
default		y	f002::1/128	Access-P	0/2	nu110	0
default		y	f002::2/128	Access-P	0/2	nu110	0
default		y	f002::3/128	Access-P	0/2	nu110	0
default		y	f002::4/128	Access-P	0/2	nu110	0
default		y	f003::2/128	Access-P	0/2	nu110	0
default		y	f004::2/128	Access-P	0/2	nu110	0
default		y	f005::2/128	Access-P	0/2	nu110	0
default		y	f006::2/128	Access-P	0/2	nu110	0
default		y	f007::2/128	Access-P	0/2	nu110	0

To specify the router context that you want to display in the output:

```
host1#show aaa route-download ipv6 routes global start Context1
```

Virtual Router	VRF	Present	Number of Routes
Context1		n	27
test		n	36

Meaning [Table 19 on page 130](#) lists the `show aaa route-download routes global` command and `show aaa route-download ipv6 routes global` command output fields.

Table 19: show aaa route-download routes global Output Fields

Field Name	Field Description
Virtual Router	Name of the virtual router used to download the routes
VRF	Name of the VRF used to download the routes
Present	Routes have been downloaded; y (yes) or n (no) indicates if the router context has been created.

Table 19: show aaa route-download routes global Output Fields (*continued*)

Field Name	Field Description
Number of Routes	Number of current downloaded routes
Prefix/Length	IP address prefix and mask information for downloaded routes
Type	Type of downloaded routes; Access-P indicates routes downloaded from the RADIUS route-download server
NextHop	IP address of the next hop
Dst/Met	Administrative distance and number of hops for the route
Tag	Tag assigned to downloaded routes
Intf	Interface type and specifier

- Related Documentation**
- *show aaa route-download routes global*
 - *show aaa route-download ipv6 routes global*

Monitoring AAA Statistics

Purpose Display authentication, authorization, and accounting statistics.

Use the optional **delta** keyword to specify that baselined statistics are to be shown.

Action To display authentication, authorization, and accounting statistics:

```
host1#show aaa statistics
```

```

          AAA Statistics
          -----
Statistic                               Count
-----
incoming initiate requests              109
incoming disconnect requests             7
outgoing grant (tunnel) responses        3
outgoing grant responses                 6
outgoing deny responses                  0
outgoing error responses                 0
outgoing Authentication requests         9
incoming Authentication responses        9
outgoing Re-Authentication requests      0
incoming Re-Authentication responses     0
outgoing Pre-Authentication requests     1
incoming Pre-Authentication responses    1
outgoing Accounting requests            120
incoming Accounting responses            120
outgoing Duplicate Acct requests         18
```

```

incoming Duplicate Acct responses      18
outgoing Broadcast Acct requests      32
incoming Broadcast Acct responses      32
outgoing Address requests              0
incoming Address responses              0

```

Meaning [Table 20 on page 132](#) lists the **show aaa statistics** command output fields.

Table 20: show aaa statistics Output Fields

Field Name	Field Description
incoming initiate requests	Number of incoming AAA requests (from other E Series applications) for user connect services
incoming disconnect requests	Number of incoming AAA requests (from other E Series applications) for user disconnect services
outgoing grant (tunnel) responses	Number of outgoing tunnel grant responses to AAA requests
outgoing grant responses	Number of outgoing grant responses to AAA requests
outgoing deny responses	Number of outgoing deny responses to AAA requests
outgoing error responses	Number of outgoing error responses to AAA requests
outgoing Authentication requests	Number of authentication requests from AAA to the authentication task
incoming Authentication responses	Number of authentication responses from the authentication task to AAA
outgoing Re-Authentication requests	Number of reauthentication requests from AAA to the authentication task
incoming Re-Authentication responses	Number of reauthentication responses from the authentication task to AAA
outgoing Pre-Authentication requests	Number of preauthentication requests from AAA to the preauthentication task
incoming Pre-Authentication responses	Number of preauthentication responses from the preauthentication task to AAA
outgoing Accounting requests	Number of accounting requests (starts, updates, stops) from AAA to the accounting task
incoming Accounting responses	Number of accounting responses (starts, updates, stops) from the accounting task to AAA
outgoing Duplicate Acct requests	Number of duplicate accounting requests (starts, updates, stops) from AAA to the accounting task

Table 20: show aaa statistics Output Fields (*continued*)

Field Name	Field Description
incoming Duplicate Acct responses	Number of duplicate accounting responses (starts, updates, stops) from the accounting task to AAA
outgoing Broadcast Acct requests	Number of broadcast accounting requests (starts, updates, stops) from AAA to the accounting task
incoming Broadcast Acct responses	Number of broadcast accounting responses (starts, updates, stops) from the accounting task to AAA
outgoing Address requests	Number of address allocation/release requests from AAA to address allocation task
incoming Address responses	Number of address allocation/release responses from the address allocation task to AAA

Related Documentation

- *show aaa statistics*

Monitoring the Number of Active Subscribers Per Port

Purpose Display the maximum number of active subscribers configured per port.

Action To display the maximum number of active subscribers configured per port:

```
host1#show aaa subscriber per-port-limit
Subscriber Port Limits
-----
Port          Limit
-----
0/2           5
0/3           2
3/2           2
```

Related Documentation

- *show aaa subscriber per-port-limit*

Monitoring the Maximum Number of Active Subscribers Per-Virtual Router

Purpose Display the maximum number of active subscribers configured per virtual router.

Action To display the maximum number of active subscribers configured per virtual router:

```
host1# show aaa subscriber per-vr-limit
subscriber limit is 0
```

Related Documentation

- *show aaa subscriber per-vr-limit*

Monitoring Session Timeouts

Purpose	Display idle and session timeouts.
Action	To display idle and session timeouts: host1#show aaa timeout idle timeout 1200 seconds monitor ingress only session timeout 3600 seconds
Related Documentation	<ul style="list-style-type: none">• <i>show aaa timeout</i>

Monitoring Interim Accounting for Users on the Virtual Router

Purpose	Display the default interval used for interim accounting for users on the virtual router. An entry of 0 indicates that the feature is disabled.
Action	To display the default interval used for interim accounting for users on the virtual router: host1:vrXyz7#show aaa user accounting interval user-acct-interval 20
Related Documentation	<ul style="list-style-type: none">• <i>show aaa user accounting interval</i>

Monitoring Virtual Router Groups Configured for AAA Broadcast Accounting

Purpose	Display the virtual router groups that are configured for AAA broadcast accounting. For additional information about the show configuration command, see <i>JunosE System Basics Configuration Guide</i> .
Action	To display the virtual router groups that are configured for AAA broadcast accounting: host1#show configuration category aaa global-attributes ! Configuration script being generated on MON JAN 10 2005 15:19:19 UTC ! Juniper Edge Routing Switch ERX1440 ! Version: 9.9.9 development-4.0 (January 7, 2005 17:26) ! Copyright (c) 1999-2004 Juniper Networks, Inc. All rights reserved. ! ! Commands displayed are limited to those available at privilege level 15 ! ! NOTE: This script represents only a subset of the full system configuration. ! The category displayed is: aaa global-attributes ! aaa accounting vr-group groupXyzCompany10 aaa virtual-router 1 vrXyzA aaa virtual-router 2 vrXyzB aaa virtual-router 3 vrXyzC aaa virtual-router 4 vrXyzD aaa accounting vr-group groupXyzCompany20 aaa virtual-router 1 vrXyzP aaa virtual-router 2 vrXyzQ aaa virtual-router 3 vrXyzR

```

aaa virtual-router 4 vrXyzS
!
hostname "host1"

```

Meaning Table 21 on page 135 lists the **show configuration category aaa global-attributes** command output fields.

Table 21: show configuration category aaa global-attributes Output Fields

Field Name	Field Description
aaa accounting vr-group	Name of virtual router groups
aaa virtual-router	Name and index number of the virtual routers that are members of the virtual router group

Related Documentation

- *show configuration*

Monitoring Configuration Information for AAA Local Authentication

Purpose Display the configuration information for AAA local authentication. You can display information for the following keywords:

- **databases**—Local user databases configured on the router
- **users**—Users configured in the local user databases
- **virtual-router**—Local user database selected by the specified virtual router for local authentication
- For additional information about the **show configuration** command, see *JunosE System Basics Configuration Guide*.

Action To display the configuration information for AAA local authentication:

```

host1#show configuration category aaa local-authentication databases
! Configuration script being generated on TUE NOV 09 2004 12:50:18 UTC
! Juniper Edge Routing Switch ERX1400
! Version: 6.1.0 (November 8, 2004 18:31)
! Copyright (c) 1999-2004 Juniper Networks, Inc. All rights reserved.
!
! Commands displayed are limited to those available at privilege level 15
!
! NOTE: This script represents only a subset of the full system configuration.
! The category displayed is: aaa local-authentication databases
!
hostname host1
aaa new-model
aaa local database default
aaa local database sval1Ldb10

```

Meaning Table 22 on page 136 lists the **show configuration category aaa local-authentication** command output fields.

Table 22: show configuration category aaa local-authentication Output Fields

Field Name	Field Description
aaa local database	Name of the local user database; the name default specifies the default local user database
aaa local select database	Local user database that the virtual router uses for local authentication
aaa local username	Unique user entry in the local user database
database	Name of the local user database for the specified username
hostname	Name of the host router
ip-address	IP address parameter for the user entry
ip-address-pool	IP address pool parameter for the user entry
operational virtual-router	Virtual router parameter for the user entry
password	Password used to authenticate the subscriber
secret	Secret used to authenticate the subscriber
virtual-router	Name of virtual router

Related Documentation

- *show configuration* category aaa local-authentication

Monitoring AAA Server Attributes

Purpose Display status of the attributes on the AAA server, including AAA accounting duplication and broadcast.

For additional information about the **show configuration** command, see *JunosE System Basics Configuration Guide*.

Action To display status of the attributes on the AAA server, including AAA accounting duplication and broadcast:

```
host1#show configuration category aaa server-attributes include-defaults
! Configuration script being generated on FRI MAY 21 2010 07:52:13 UTC
! Juniper Edge Routing Switch ERX1440
! Version: 11.2.0 beta-1.1 [BuildId 12073] (April 22, 2010 11:46)
! Copyright (c) 1999-2010 Juniper Networks, Inc. All rights reserved.
!
! Commands displayed are limited to those available at privilege level 15
!
! NOTE: This script represents only a subset of the full system configuration.
```



```
! The category displayed is: aaa server-attributes
!
virtual-router default
aaa accounting duplication lac
aaa accounting broadcast group1
aaa duplicate-address-check enable
aaa accounting acct-stop on-aaa-failure enable
aaa accounting acct-stop on-access-deny disable
aaa subscriber limit per-vr 0
aaa intf-desc-format include sub-intf enable
aaa intf-desc-format include adapter enable
aaa accounting immediate-update disable
no aaa ipv6-nd-ra-prefix framed-ipv6-prefix
no aaa dhcpv6-delegated-prefix delegated-ipv6-prefix
aaa duplicate-prefix-check disable
!
! =====
!
virtual-router lac
no aaa accounting duplication
no aaa accounting broadcast
aaa duplicate-address-check enable
aaa accounting acct-stop on-aaa-failure enable
aaa accounting acct-stop on-access-deny disable
aaa subscriber limit per-vr 0
aaa intf-desc-format include sub-intf enable
aaa intf-desc-format include adapter enable
aaa accounting immediate-update disable
no aaa ipv6-nd-ra-prefix framed-ipv6-prefix
no aaa dhcpv6-delegated-prefix delegated-ipv6-prefix
aaa duplicate-prefix-check disable
!
! =====
!
virtual-router isp
no aaa accounting duplication
no aaa accounting broadcast
aaa duplicate-address-check enable
aaa accounting acct-stop on-aaa-failure enable
aaa accounting acct-stop on-access-deny disable
aaa subscriber limit per-vr 0
aaa intf-desc-format include sub-intf enable
aaa intf-desc-format include adapter enable
aaa accounting immediate-update disable
no aaa ipv6-nd-ra-prefix framed-ipv6-prefix
no aaa dhcpv6-delegated-prefix delegated-ipv6-prefix
aaa duplicate-prefix-check disable
```

Meaning [Table 23 on page 137](#) lists the **show configuration category aaa server-attributes include-defaults** command output fields.

Table 23: show configuration category aaa server-attributes include-defaults Output Fields

Field Name	Field Description
virtual router	Name of the virtual router
aaa accounting duplication	Virtual router used for duplicate accounting

Table 23: show configuration category aaa server-attributes include-defaults Output Fields (*continued*)

Field Name	Field Description
aaa accounting broadcast	Virtual router group used for broadcast accounting
aaa duplicate-address-check	Enabled, disabled
aaa accounting acct-stop on-aaa-failure	Enabled, disabled
aaa accounting acct-stop on-access-deny	Enabled, disabled
aaa subscriber limit per-vr	Enabled, disabled
aaa intf-desc-format include sub-intf	Enabled, disabled
aaa intf-desc-format include adapter	Enabled, disabled
aaa accounting immediate-update	Enabled, disabled
aaa ipv6-nd-ra-prefix framed-ipv6-prefix	Framed-IPv6-Prefix RADIUS attribute used for IPv6 Neighbor Discovery router advertisements
aaa dhcpv6-delegated-prefix delegated-ipv6-prefix	Delegated-IPv6-Prefix RADIUS attribute used for DHCPv6 prefix delegation
aaa duplicate-prefix-check	Enabled, disabled

Related Documentation

- *show configuration*

Monitoring the COPS Layer Over SRC Connection

Purpose Display information about the COPS layer over which the SRC connection is made.

Action To display information about the COPS layer over which the SRC connection is made:

```
host1#show cops info
```

```
General Cops Information:
```

```
Sessions Created: 1
Sessions Deleted: 0
Current Sessions: 1
Bytes Received: 680
Packets Received: 17
Bytes Sent: 692
Packets Sent: 21
Keep Alive Received: 12
```

Keep Alive Sent: 12

Session Information

Remote Ip Address: 10.10.0.223
 Remote TCP Port: 4001
 Client Type: 16384
 Bytes Received: 2224
 Packets Received: 5
 Bytes Sent: 596
 Packets Sent: 9
 REQ Sent: 4
 DEC Rcv: 4
 RPT Sent: 4
 DRQ Sent: 0
 SSQ Rcv: 0
 OPN Sent: 1
 CAT Rcv: 1
 CC Sent: 0
 CC Rcv: 0
 SSC Sent: 0

Meaning [Table 24 on page 139](#) lists the **show cops info** command output fields.

Table 24: show cops info Output Fields

Field Name	Field Description
Session Created	Number of COPS sessions created
Sessions Deleted	Number of COPS sessions deleted
Current Sessions	Number of current COPS sessions
Bytes Received	Number of bytes received on all COPS sessions
Packets Received	Number of packets received on all COPS sessions
Bytes Sent	Number of bytes transmitted on all COPS sessions
Packets Sent	Number of packets transmitted on all COPS sessions
Keep Alive Received	Number of COPS keepalive messages received
Keep Alive Sent	Number of COPS keepalive messages <i>sent</i>
Remote IP Address	IP address of the remote peer
Remote TCP Port	TCP port number of the remote peer
Client Type	Type of client for the session. For this release the client type must be 16640 (SRC client).
Bytes Received	Number of bytes received for this COPS session
Packets Received	Number of packets received for this COPS session

Table 24: show cops info Output Fields (*continued*)

Field Name	Field Description
Bytes Sent	Number of bytes sent on this COPS session
Packets Sent	Number of packets sent on this COPS session
REQ Sent	Number of Request packets sent on this COPS session
DEC Rcv	Number of Decision packets received on this COPS session
RPT Sent	Number of Report packets sent on this COPS session
DRQ Sent	Number of Delete Requests sent on this COPS session
SSQ Rcv	Number of Synch Requests received on this COPS session
OPN Sent	Number of Open messages sent on this COPS session
CAT Rcv	Number of Client Accepts packets received on this COPS session
CC Sent	Number of Client Closes packets sent on this COPS session
CC Rcv	Number of Client Closes packets received on this COPS session
SSC Sent	Number of Sync Complete packets sent on this COPS session

Related Documentation

- *show cops info*

Monitoring Statistics About the COPS Layer

Purpose Display statistics about the COPS layer over which the SRC connection is made.

Action To display statistics about the COPS layer:

```
host1#show cops statistics
General Cops Information:
  Sessions Created: 0
  Sessions Deleted: 0
  Current Sessions: 0
  Bytes Received: 1108
  Packets Received: 12
  Bytes Sent: 1572
  Packets Sent: 18
  Keep Alive Received: 2
```

```

Keep Alive Sent: 2
Session Information:
Client Type: 24754
Bytes Received: 2539032
Packets Received: 20388
Bytes Sent: 4386648
Packets Sent: 51337
REQ Sent: 21203
DEC Rcv: 20388
RPT Sent: 20391
DRQ Sent: 9743
SSQ Rcv: 0
OPN Sent: 0
CAT Rcv: 0
CC Sent: 0
CC Rcv: 0
SSC Sent: 0

```

Meaning [Table 25 on page 141](#) lists the **show cops statistics** command output fields.

Table 25: show cops statistics Output Fields

Field Name	Field Description
Session Created	Number of COPS sessions created
Sessions Deleted	Number of COPS sessions deleted
Current Sessions	Number of current COPS sessions
Bytes Received	Number of bytes received on all COPS sessions
Packets Received	Number of packets received on all COPS sessions
Bytes Sent	Number of bytes transmitted on all COPS sessions
Packets Sent	Number of packets transmitted on all COPS sessions
Keep Alive Received	Number of COPS keepalive messages received
Keep Alive Sent	Number of COPS keepalive messages <i>sent</i>
Client Type	Type of client for the session
Bytes Received	Number of bytes received for this COPS session
Packets Received	Number of packets received for this COPS session
Bytes Sent	Number of bytes sent on this COPS session
Packets Sent	Number of packets sent on this COPS session
REQ Sent	Number of Request packets sent on this COPS session

Table 25: show cops statistics Output Fields (*continued*)

Field Name	Field Description
DEC Rcv	Number of Decision packets received on this COPS session
RPT Sent	Number of Report packets sent on this COPS session
DRQ Sent	Number of Delete Requests sent on this COPS session
SSQ Rcv	Number of Synch Requests received on this COPS session
OPN Sent	Number of Open messages sent on this COPS session
CAT Rcv	Number of Client Accepts packets received on this COPS session
CC Sent	Number of Client Closes packets sent on this COPS session
CC Rcv	Number of Client Closes packets received on this COPS session
SSC Sent	Number of Sync Complete packets sent on this COPS session

Related Documentation

- *show cops statistics*

Monitoring Local Address Pool Aliases

Purpose Display information about aliases for the local address pools configured on your router. If you do not specify a particular alias, the router displays all aliases.

Action To display information about local address pool aliases:

```
host1#show ip local alias
```

```
Alias      Pool
-----
alias1     poolA
alias2     poolB
alias3     poolC
poolA      poolD
poolB      poolD
poolC      poolD
```

Meaning [Table 26 on page 143](#) lists the show **ip local alias** command output fields.

Table 26: show ip local alias Output Fields

Field Name	Field Description
Alias	Name of alias for the local address pool
Pool	Name of the local address pool

Related Documentation

- *show ip local alias*

Monitoring Local Address Pools

Purpose Display information about the local address pools configured on your router. If you do not specify the name of a local address pool, the router displays all local address pools.

Action To display information about local address pools:

```
host1#show ip local pool
```

```

Pool      High    Abated
-----  Thresh  Thresh  Trap   Group
-----  -----
poolA      85      75      N

```

Aliases

alias1

```

      Begin      End      Free    In
      -----    -----    ----    --
10.1.1.1  10.1.1.10    10      0
10.1.2.1  10.1.2.10    10      0
10.1.3.1  10.1.3.10    10      0

```

```

Pool      High    Abated
-----  Thresh  Thresh  Trap   Group
-----  -----
poolB      85      75      N

```

Aliases

alias2

```

      Begin      End      Free    In
      -----    -----    ----    --
10.2.1.1  10.2.1.10    10      0
10.2.2.1  10.2.2.10    10      0

```

```

Pool      High    Abated
-----  Thresh  Thresh  Trap   Group
-----  -----
poolC      85      75      N

```

Aliases

alias3

```

      Begin      End      Free    In
      -----    -----    ----    --

```

```

-----
10.3.1.1  10.3.1.10  10  0
Pool      High    Abated
Thresh    Thresh    Trap  Group
-----
poolD      85      75    N
Aliases
-----
poolA
poolB
poolC
Begin      End      Free    In
-----    -----
10.4.1.1  10.4.1.255  255    0

```

Meaning [Table 27 on page 144](#) lists the **show ip local pool** command output fields.

Table 27: show ip local pool Output Fields

Field Name	Field Description
Pool	User-specified name of the address pool
High Thresh	High utilization threshold value
Abated Thresh	Abated utilization threshold value
Trap	Enable SNMP pool utilization traps: Y (yes) or N (no)
Aliases	Aliases for the local address pool
Begin	Starting IP address
End	Ending IP address
Free	Number of addresses available for use
In Use	Number of addresses currently in use

Related Documentation

- *show ip local pool*

Monitoring Local Address Pool Statistics

Purpose Display local address pool statistics. Use the optional **delta** keyword to specify that baselined statistics are to be shown.

Action To display local address pool statistics:

```

host1#show ip local pool statistics
Local Address Pool Statistics

```


Statistic	Values
Requests denied (pool exhaustion)	0

Related Documentation

- show ip local pool*

Monitoring Shared Local Address Pools

Purpose Display the shared local address pool configurations.

Action To display shared local address pool configuration information:

```
host1#show ip local shared-pool
```

Shared Pool	In Use	Dhcp Pool
shared_poolA	253	dhcp_pool_25
shared_poolB	83	dhcp_pool_25
shared_poolC	99	dhcp_pool_17

Meaning [Table 28 on page 145](#) lists the **show ip local shared-pool** command output fields.

Table 28: show ip local shared-pool Output Fields

Field Name	Field Description
Shared Pool	Name of the shared local address pool
In Use	Number of addresses allocated
Dhcp Pool	Name of the DHCP address pool

Related Documentation

- show ip local shared-pool*

Monitoring the Routing Table

Purpose Display the current state of the routing table, including routes not used for forwarding. An Access-P entry in the Type column of the output indicates routes that are downloaded by the RADIUS route-download server.

Action To display information in the routing table:

```
host1#show ip route
```

Protocol/Route type codes:
I1- ISIS level 1, I2- ISIS level2,
I- route type intra, IA- route type inter, E- route type external,
i- metric type internal, e- metric type external,
P- periodic download, O- OSPF, E1- external type 1, E2- external type2,
N1- NSSA external type1, N2- NSSA external type2
L- MPLS label, V- VRF, *- via indirect next-hop

Prefix/Length	Type	Next Hop	Dst/Met	Interface
-----	-----	-----	-----	-----

0.0.0.0/0	Static	10.13.10.1	1/0	FastEthernet6/0/0
192.168.10.0/23	Connect	10.13.10.187	0/0	FastEthernet6/0/0
192.168.21.21/32	Access-P	255.255.255.255	254/2	null0
192.168.22.22/32	Access-P	255.255.255.255	254/2	null0
192.168.23.23/32	Access-P	255.255.255.255	254/2	null0
192.168.24.24/32	Access-P	255.255.255.255	254/2	null0

Meaning Refer to the description of the **show ip route** command in *JunosE IP, IPv6, and IGP Configuration Guide* for additional information about the **show ip route** command.

Related Documentation

- *show ip route*

Monitoring the B-RAS License

Purpose Display the B-RAS license.

Action To display the B-RAS license:

```
host1#show license b-ras
K4bZ16Lr
```

Related Documentation

- *show license b-ras*

Monitoring the RADIUS Server Algorithm

Purpose Display information about the currently configured RADIUS server algorithm.

Action To display the RADIUS server algorithm:

```
host1#show radius algorithm
direct
```

Related Documentation

- *show radius algorithm*

Monitoring RADIUS Override Settings

Purpose Display the current RADIUS override settings.

Action To display the RADIUS override settings:

```
host1:vrXyz7#show radius override
nas-ip-addr: nas-ip-addr
nas-info:    from authentication virtual router
```

Meaning [Table 29 on page 147](#) lists the **show radius override** command output fields.

Table 29: show radius override Output Fields

Field Name	Field Description
nas-ip-addr	Either the NAS-IP-Address [4] attribute is used, or it is overridden with the Tunnel-Client-Endpoint [66] attribute.
nas-info	Either the NAS-IP-Address [4] and NAS-Identifier [32] attributes of the virtual router generating the accounting information are used, or they are overridden with the respective attributes of the authentication virtual router.

Related Documentation

- *show radius override*

Monitoring the RADIUS Rollover Configuration

Purpose Display the configuration of the RADIUS rollover-on-reject feature.

Action To display the RADIUS rollover configuration:

```
host1#show radius rollover-on-reject
rollover-on-reject enabled
```

Meaning RADIUS rollover-on-reject is enabled.

Related Documentation

- *show radius rollover-on-reject*

Monitoring RADIUS Server Information

Purpose Display RADIUS server information.

Use with the optional **accounting**, **authentication**, **dynamic-request**, **route-download**, or **pre-authentication** keywords to limit output to the specific type of server.

Action To display RADIUS server configuration information:

```
host1#show radius servers
```

RADIUS Authentication Configuration							
IP Address	Udp Port	Retry Count	Timeout	Maximum Sessions	Dead Time	Secret	Status
172.28.30.117	1812	3	3	255	30	radius	dead
172.28.30.118	1812	3	3	255	30	radius	active
172.28.30.119	1812	3	3	255	30	radius	alive

RADIUS Accounting Configuration							
IP Address	Udp Port	Retry Count	Timeout	Maximum Sessions	Dead Time	Secret	Status
172.28.30.117	1813	3	3	255	30	radius	dead

```

172.28.30.118 1813 3 3 255 30 radius active
172.28.30.119 1813 3 3 255 30 radius alive

```

RADIUS Pre-Authentication Configuration

IP Address	Udp Port	Retry Count	Timeout	Maximum Sessions	Dead Time	Secret	Status
172.28.30.117	1812	3	3	255	30	radius	dead
172.28.30.118	1812	3	3	255	30	radius	active
172.28.30.119	1812	3	3	255	30	radius	alive

RADIUS Route-Download Configuration

IP Address	Udp Port	Retry Count	Timeout	Maximum Sessions	Dead Time	Secret	Status
192.168.30.16	1812	3	3	255	30	radius	dead
192.168.30.17	1812	3	3	255	30	radius	active
192.168.30.18	1812	3	3	255	30	radius	alive

Meaning If a RADIUS server was never configured on the virtual router, the command displays the following message:

```

host1#show radius servers
no radius servers configured

```

If a RADIUS server was configured previously and then removed on the virtual router, the command displays the following information:

```

host1#show radius servers
RADIUS Authentication Configuration
-----
IP Address  Udp Port  Retry Count  Timeout  Maximum Sessions  Dead Time  Secret  Status
-----
RADIUS Accounting Configuration
-----
IP Address  Udp Port  Retry Count  Timeout  Maximum Sessions  Dead Time  Secret  Status
-----
RADIUS Pre-Authentication Configuration
-----
IP Address  Udp Port  Retry Count  Timeout  Maximum Sessions  Dead Time  Secret  Status
-----
RADIUS Route-Download Configuration
-----
IP Address  Udp Port  Retry Count  Timeout  Maximum Sessions  Dead Time  Secret  Status
-----

```

Table 30 on page 148 lists the **show radius servers** command output fields.

Table 30: show radius servers Output Fields

Field Name	Field Description
IP Address	IP address of RADIUS server

Table 30: show radius servers Output Fields (*continued*)

Field Name	Field Description
Udp Port	Number of the UDP port of the RADIUS server
Retry Count	Maximum number of times that the router retransmits a RADIUS packet to the RADIUS server
Timeout	Interval (in seconds) before the router retransmits a RADIUS packet to the RADIUS server
Maximum Sessions	Number of outstanding requests to the RADIUS server
Dead Time	Amount of time to remove the authentication server or accounting server from the available list when a timeout occurs
Secret	Configured authentication server or accounting server secret
Status	Status of the configured RADIUS server: <ul style="list-style-type: none"> • dead-The status displayed if the server does not respond within the configured number of retransmit counts, and if Dead Time is configured to a non-zero value. • active-The status displayed of the earliest configured, non-dead server if the server is accessed using the direct algorithm. The status displayed of all non-dead servers if the server is accessed using the round-robin algorithm. • alive-The status displayed of all non-dead servers except the earliest configured non-dead server, if the server is accessed using the direct algorithm. The status of none of the servers if the server is accessed using the round-robin algorithm.

Related Documentation

- *show radius servers*

Monitoring RADIUS Services Statistics

Purpose Use to display statistics for RADIUS services.

Use with the optional **accounting**, **authentication**, **dynamic-request**, **route-download**, or **pre-authentication** keywords to limit output to the specific type of statistics. Use the optional **delta** keyword to specify that baselined statistics are to be shown.

Action To display RADIUS authentication and accounting statistics:

```
host1#show radius statistics
      RADIUS Authentication Statistics
-----
Statistic      10.10.121.128
-----
```

UDP Port	1812
Round Trip Time	0
Access Requests	0
Rollover Requests	0
Retransmissions	0
Access Accepts	0
Access Rejects	0
Access Challenges	0
Malformed Responses	0
Bad Authenticators	0
Requests Pending	0
Request Timeouts	0
Unknown Responses	0
Packets Dropped	0

RADIUS Accounting Statistics

Statistic	10.10.121.128
UDP Port	1646
Round Trip Time	2
Requests	1
Start Requests	1
Interim Requests	0
Stop Requests	0
Reject Requests	0
Rollover Requests	0
Retransmissions	3
Responses	1
Start Responses	1
Interim Responses	0
Stop Responses	0
Reject Responses	0
Malformed Responses	0
Bad Authenticators	0
Requests Pending	0
Request Timeouts	3
Unknown Responses	0
Packets Dropped	0

To display RADIUS pre-authentication statistics:

```
host1#show radius pre-authentication statistics
```

RADIUS Pre-Authentication Statistics

Statistic	172.28.30.117
UDP Port	1812
Round Trip Time	0
Access Requests	2809
Rollover Requests	0
Retransmissions	56
Access Accepts	2809
Access Rejects	0
Access Challenges	0
Malformed Responses	0
Bad Authenticators	0
Requests Pending	0
Request Timeouts	72
Unknown Responses	0
Packets Dropped	2

To display RADIUS route-download statistics:

```
host1#show radius route-download statistics
```

```

RADIUS Route-Download Statistics
-----
Statistic          192.168.30.16
-----
UDP Port           1812
Round Trip Time    0
Access Requests    1613
Rollover Requests  0
Retransmissions    6
Access Accepts     1612
Access Rejects     1
Access Challenges  0
Malformed Responses 0
Bad Authenticators 0
Requests Pending   0
Request Timeouts   6
Unknown Responses  0
Packets Dropped    5

```

Meaning [Table 31 on page 151](#) lists the **show radius statistics** command output fields.



NOTE: All descriptions apply to the primary, secondary, and tertiary RADIUS authentication and accounting servers.

Table 31: show radius statistics Output Fields

Field Name	Field Description
UDP Port	Number of the UDP port of a RADIUS server
Round Trip Time	Time interval between a request sent to the server and its response received from the server. The interval value is displayed in centiseconds (one-hundredth of a second).
Access Requests	Number of access requests sent to server
Rollover Requests	Number of requests coming into server as a result of the previous server timing out
Retransmissions	Number of retransmissions
Access Accepts	Number of Access-Accepts received from the server
Access Rejects	Number of Access-Rejects received from the server
Access Challenges	Number of access challenges received from the server

Table 31: show radius statistics Output Fields (*continued*)

Field Name	Field Description
Malformed Responses	Number of responses with attributes having an invalid length or unexpected attributes (such as two attributes when the response is required to have at most one)
Bad Authenticators	Number of responses in which the authenticator is incorrect for the matching request. This can occur if the RADIUS secret for the client and server does not match.
Requests Pending	Number of requests waiting for a response
Request Timeouts	Number of requests that timed out
Unknown Responses	Number of unknown responses. The RADIUS response type in the header is invalid or unsupported.
Packets Dropped	Number of packets dropped either because they are too short or the E Series router receives a response for which there is no corresponding request. For example, if the router sends a request and the request 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.
Requests	Total number of accounting requests sent, which is the combined total of Start Requests, Interim Requests, Stop Requests, and Reject Requests
Start Requests	Number of accounting start requests sent; includes Acct-On, Acct-Start, Acct-Link-State, and Acct-Tunnel-Start requests
Interim Requests	Number of interim accounting requests
Stop Requests	Number of accounting stop requests sent; includes Acct-Off, Acct-Stop, Acct-Link-Stop, and Acct-Tunnel-Stop requests
Reject Requests	Number of accounting reject requests sent; includes Acct-Link-Reject and Acct-Tunnel-Reject requests
Responses	Number of accounting responses received from the server
Start Responses	Number of accounting start responses received; includes Acct-On, Acct-Start, Acct-Link-Start, and Acct-Tunnel-Start responses
Interim Responses	Number of interim accounting responses

Table 31: show radius statistics Output Fields (*continued*)

Field Name	Field Description
Stop Responses	Number of accounting stop responses received; includes Acct-Off, Acct-Stop, Acct-Link-Stop, and Acct-Tunnel-Stop responses
Reject Responses	Number of accounting reject responses received; includes Acct-Link-Reject and Acct-Tunnel-Reject responses

Related Documentation

- *show radius statistics*

Monitoring RADIUS SNMP Traps

Purpose Display the configuration of RADIUS SNMP traps.

Action To display RADIUS SNMP traps configuration information:

```
host1#show radius trap
trap for auth-server-not-responding enabled
trap for no-auth-server-responding disabled
trap for auth-server-responding enabled
trap for acct-server-not-responding enabled
trap for no-acct-server-responding disabled
trap for acct-server-responding disabled
```

Meaning A list of the configured RADIUS-related SNMP traps is displayed.

Related Documentation

- *show radius trap*

Monitoring RADIUS Accounting for L2TP Tunnels

Purpose Display the status for RADIUS accounting for L2TP tunnels.

Action To display RADIUS accounting for L2TP tunnels:

```
host1#show radius tunnel-accounting
disabled
```

Meaning RADIUS accounting is either enabled or disabled.

Related Documentation

- *show radius tunnel-accounting*

Monitoring RADIUS UDP Checksums

Purpose Display information about UDP checksums.

Action To display the status of RADIUS UDP checksums:

```
host1#show radius udp-checksum
enab1ed
```

Meaning RADIUS checksums status is either enabled or disabled.

Related Documentation • *show radius udp-checksum*

Monitoring RADIUS Server IP Addresses

Purpose Display the IP address of the RADIUS servers.

Action To display the RADIUS server IP address:

```
host1#show radius update-source-address
192.168.1.228
```

Related Documentation • *show radius update-source-addr*

Monitoring the RADIUS Attribute Used for IPv6 Neighbor Discovery Router Advertisements

Purpose Display the RADIUS attribute used for IPv6 Neighbor Discovery router advertisements.

Action To display the RADIUS attribute used for IPv6 Neighbor Discovery router advertisements:

```
host1#show aaa ipv6-nd-ra-prefix
IPv6 ND RA Prefix      : IPv6-NdRa-Prefix (Juniper VSA)
```

Related Documentation • *show aaa ipv6-nd-ra-prefix*

Monitoring the RADIUS Attribute Used for DHCPv6 Prefix Delegation

Purpose Display the RADIUS attribute used for DHCPv6 Prefix Delegation.

Action To display the RADIUS attribute used for DHCPv6 Prefix Delegation:

```
host1#show aaa dhcpv6-delegated-prefix
DHCPv6 Delegated Prefix : Framed-IPv6-Prefix
```

Related Documentation • *show aaa dhcpv6-delegated-prefix*

Monitoring Duplicate IPv6 Prefixes

Purpose Display whether the ability to detect duplicates of IPv6 Neighbor Discovery router advertisement prefixes and DHCPv6 delegated prefixes is enabled.

Action To check whether duplicate IPv6 prefix detection capability is enabled:

```
host1#show aaa duplicate-prefix-check
enabled
```

Related Documentation

- [show aaa duplicate-prefix-check](#)

Monitoring Duplicate IPv6 Prefixes in the AAA User Profile Database

Purpose Display whether the ability to detect duplicates of IPv6 Neighbor Discovery router advertisement prefixes and DHCPv6 delegated prefixes, in the AAA userProfile database, is enabled.

Action To check whether enhanced duplicate IPv6 prefix detection capability is enabled:

```
host1#show aaa duplicate-prefix-check-extension
enabled
```

Related Documentation

- [Duplicate IPv6 Prefix Detection in the AAA User Profile Database Overview on page 42](#)
- [Configuring Detection of Duplicate IPv6 Prefixes in the AAA User Profile Database on page 97](#)
- [show aaa duplicate-prefix-check-extension](#)

Monitoring the SRC Client Connection Status

Purpose Display the current status of the SRC client connection to the service activation engines (SAEs). The command output refers to the SRC client by its former name: SSC client.

Action To display the status of the SRC client connection:

```
host1#show sssc info
The SSC Client is connected to: 10.212.10.17:3288
The SSC Client configured protocols : IP(v4), DHCP(v4), IP(v6), DHCP(v6)
The SSC Client configured servers are:
  Primary: 10.212.10.17:3288
  Secondary: 0.0.0.0:0
  Tertiary: 0.0.0.0:0
Local Source: not configured, Local Source Address: 0.0.0.0
The configured transport router is: default
The configured retry timer is (seconds): 90
The configured connectivity timer is (seconds): 0
The configured update-policy-request is: Disabled
The connection state is: open
SSC Client Statistics:
DEC messages received          12
Config decisions received      0
Accounting decisions received  0
Decisions failed               2
Error decisions received       0
RPT messages sent             12
OPN message sent               1
CAT messages received          1
CC messages sent               0
CC messages received           0
```

Interface creations	6
Interface deletions	0
Interfaces active	6
Interface transitions	0
SSQ messages received	1
COPS sessions opened	1
SSC messages sent	1
Internal errors	0
Comm errors	0
Discover notifications	0
Discovers active	0
Discover transitions	0
Discover creations	0
Discovers deletions	0
Address active	0
Address transitions	0
Address creations	0
Address deletions	0

Meaning Table 32 on page 156 lists the **show sssc info** command output fields.

Table 32: show sssc info Output Fields

Field Name	Field Description
The SSC Client is connected to	IP address to which the SSC client is connected
The SSC Client configured protocols	Protocols that are enabled on the virtual router for policy and QoS management by the SRC software
The SSC Client configured servers are	IP addresses of the primary, secondary, and tertiary SAEs
Local Source	Fixed source interface for the TCP/COPS connection
Local Source Address	Fixed source address for the TCP/COPS connection
The configured transport router is	Router on which the TCP/COPS connection is established
The configured retry timer is (seconds)	Delay period during which the client waits for a response from the SAE before submitting the request again
The configured connectivity timer is (seconds)	Time period until which the rebooted SRC client waits for the primary SRC server to establish a connection with the SRC client
The configured update-policy-request is	Whether the router or the SRC client retrieves DSL line-rate parameters, whenever the values change after connection establishment, from ANCP and transfers the details to the COPS server with other COPS messages: enabled or disabled
The connection state is	Current state of the TCP/COPS connection

Table 32: show ssc info Output Fields (*continued*)

Field Name	Field Description
SSC Client Statistics	<p>Statistics about the connection between the SRC client and the SAE:</p> <ul style="list-style-type: none"> • DEC messages received—Number of COPS decision messages received on the router from the SRC client • Config decisions received—Number of configuration decision messages received from the SRC client • Accounting decisions received—Number of accounting decision messages received from the SRC client • Decisions failed—Number of decision messages that failed on the router during processing • Error decisions received—Number of error decisions received from the SRC client • RPT messages sent—Number of COPS RPT messages sent from the router to the SRC client • OPN message sent—Number of COPS OPN messages sent to the SRC client • CAT messages received—Number of COPS CAT messages received from the SRC client • CC messages sent—Number of COPS CC messages sent to the SRC client • CC messages received—Number of COPS CC messages received from the SRC client • Interface creations—Number of interfaces that are created • Interface deletions—Number of Interfaces that are deleted • Interfaces active—Number of interfaces that are currently active • Interface transitions—Number of interfaces that transitioned from the up state to the down state • SSQ messages received—Number of SSQ messages received from the SRC client • COPS sessions opened—Number of COPS sessions opened • SSC messages sent—Number of SSC messages sent to the SRC client • Internal errors—Internal processing errors • Comm errors—Communication errors associated with the SRC client • Discover notifications—Number of discover notifications that the SSC client received from the DHCP server • Discovers active—Number of discover messages that are currently active • Discover transitions—Number of discover messages that transitioned from the up state to the down state • Discover creations—Number of DHCP discover messages that are created • Discovers deletions—Number of DHCP discover messages that are deleted • Address active—Number of DHCP subscriber addresses that are currently active • Address transitions—Number of DHCP subscriber addresses that transitioned from the active state to the inactive state • Address creations—Number of DHCP addresses that are created • Address deletions—Number of DHCP addresses that are deleted

- Related Documentation**
- [Configuring the SRC Client on page 97](#)
 - *show sssc info*

Monitoring SRC Client Connection Statistics

Purpose Display statistics about connection between the SRC client and SAE. The command output refers to the SRC client by its former name, SSC client.

Action To display statistics for the SRC client connection:

```
host1#show sssc statistics
```

```
SSC Client Statistics:
```

```
Policy Commands received      0
Policy Commands(List)         0
Policy Commands(Acct)         0
Bad Policy Cmds received      0
Error Policy Cmds received    0
Policy Reports sent           3
Connection attempts           7
Connection Open requests      7
Connection Open completed     0
Connection Closed sent        0
Connection Closed remotely    5
Create Interfaces sent         0
Delete Interfaces sent         3
Active IP Interfaces           3282
IP Interface Transitions       3281
Synchronizes received         0
Synchronizes rcvd & dropped    0
Synchronize Complete sent     2
Internal Errors                0
Communication Errors           0
Discovers Seen                 15263
Active Discovers               4911
Discover Transitions           20704
Discover Creates Sent          15263
Discover Deletes Sent          10352
Active Addresses               3274
Address Transitions            3280
Create Addresses Sent          3277
Delete Addresses Sent          3
```

Meaning [Table 33 on page 158](#) lists the **show sssc statistics** command output fields.

Table 33: show sssc statistics Output Fields

Field Name	Field Description
Policy Commands received	Number of policy commands received on the SRC client connection
Policy Commands(List)	Number of Policy Commands with subtype List
Policy Commands(Acct)	Number of Policy Commands with subtype Accounting

Table 33: show sssc statistics Output Fields (*continued*)

Field Name	Field Description
Bad Policy Cmds received	Number of Policy Commands received with bad policies
Error Policy Cmds received	Number of Policy Commands received with errors
Policy Reports sent	Number of Policy Reports sent
Connection Open requests	Number of connections the SRC client has tried to open with a remote SAE
Connection Open completed	Number of connections successfully open to the SAE
Connection Closed sent	Number of connections the SRC client has closed
Connection Closed remotely	Number of connections that were closed by the remote SAE
Create Interfaces sent	Number of create interface indications sent to the SAE
Delete Interfaces sent	Number of delete interface indications sent to the SAE
Active IP Interfaces	Current number of active IP interfaces the SRC client is aware of
IP Interface Transitions	Number of IP interface transitions logged by the SRC client
Synchronizes received	Number of synchronization requests the SRC client received from the SAE
Synchronize Complete sent	Number of synchronization complete indications sent
Internal Errors	Number of internal errors
Communication Errors	Number of errors with lower-layer communications (such as socket errors)

Related Documentation

- *show sssc statistics*

Monitoring the SRC Client Version Number

Purpose Display the SRC client (formerly SDX client) version number.

Action To display the SRC client version number:

```
host1#show sssc version
The SSC Client version is: 4.0
```

Related Documentation

- *show sssc version*

Monitoring the SRC Option

Purpose Displays information about SRC client options for the virtual router.

Action To display the SRC option:

```
host1#show sssc options
The SSC Client options for vr default:
  generate-nas-port-id: disabled
    Send-Calling-Station-Id with radius-default-value: enabled
  Send-Calling-Station-Id with radius-overridden-value: disabled
  send-lac-nas-ip: enabled
  send-lac-nas-port: enabled
  send-local-qos-profile-config: disabled
  user-ip-mask-override: disabled
```

Meaning [Table 34 on page 160](#) lists the **show sssc option** command output fields.

Table 34: show sssc option Output Fields

Field Name	Field Description
generate-nas-port-id	If enabled, the LNS side NAS-Port information is sent to the PDP for a virtual router
Send-Calling-Station-Id with radius-default-value	If enabled, the default calling station ID is sent to the PDP for a virtual router, irrespective of whether the ID is included or excluded from the Access-Request and Acct-Start messages sent from the RADIUS client to the server
Send-Calling-Station-Id with radius-overridden-value	If enabled, the overridden calling station ID is sent to the PDP for a virtual router, irrespective of whether the ID is included or excluded from the Access-Request and Acct-Start messages sent from the RADIUS client to the server
send-calling-station-id	If enabled, the calling station ID is sent to the PDP for a virtual router
send-lac-nas-ip	If enabled, the LAC side NAS-IP address information is sent to the PDP for a virtual router
send-lac-nas-port	If enabled, the LAC side NAS-Port information is sent to the PDP for a virtual router
send-local-qos-profile-config	If enabled, the local QoS profile attachment information is sent to the PDP for a virtual router

Table 34: show sssc option Output Fields (*continued*)

Field Name	Field Description
user-ip-mask-override	If enabled, the user IP address mask is sent to the PDP for a virtual router

- Related Documentation**
- *show sssc options*
 - *sssc option*

Monitoring Subscriber Information

- Purpose** Display active subscribers on the router. If you specify a username, the router displays only the users that match the username. When you issue the **show subscribers** command in the default virtual router, all users are displayed. When you issue the **show subscribers** command in a nondefault virtual router, only those users attached to that virtual router are displayed. The following list describes keywords that you can issue with the **show subscribers** command:
- You can specify the **domain**, **interface**, **port**, **profile**, **slot**, **username**, or **virtual-router** keyword on all routers to filter the results. If you do not specify a keyword, all active users are displayed.
 - When you use the **interface** keyword to display detailed subscriber information by interface, you must also specify the **atm**, **ethernet**, or **lag** keyword, an interface specifier, and, optionally, a subinterface specifier.
 - When you use the **username** keyword to display detailed subscriber information using the subscriber username, you must specify the complete username with the domain name (if available) regardless of the status of the strip domain feature on a virtual router or AAA domain map.
 - If you specify the **lag** keyword, the output displays active subscribers for the specified LAG interface. By default, the **aaa intf-desc-format include sub-intf enable** command includes the subinterface and S-VLAN ID in the LAG interface ID. Use the **aaa intf-desc-format include sub-intf disable** command to exclude the subinterface and S-VLAN ID from the LAG interface ID.
 - The output displayed in the Interface field depends on the configuration of two commands at the time the subscriber logs in: **aaa intf-desc-format include sub-intf** and **aaa intf-desc-format include adapter** (for the E120 and E320 Broadband Services routers).
 - When you issue the **aaa intf-desc-format include sub-intf disable** command, the subinterface is stripped from the subscriber's interface field at login and is not displayed in the output. In the default state, or when you issue the **aaa intf-desc-format include sub-intf enable** command, the subinterface is included in the subscriber's interface field at login and is displayed in the output.
 - When you issue the **aaa intf-desc-format include adapter disable** command, the adapter is stripped from the subscriber's interface field at login and is not displayed

in the output. In the default state, or when you issue the **aaa intf-desc-format include adapter enable** command, the adapter is included in the subscriber's interface field at login and is displayed in the output.

- Even when the subinterface has been stripped from the subscriber's interface field, you can still include the subinterface specifier in the **show subscribers interface** command. Even though the subinterface itself is not displayed, only subscribers on the specified subinterface are displayed.
- The preceding considerations do not apply when you issue the **summary** keyword. The output displayed in the Interface field of summary versions is not affected by the state of either the **aaa intf-desc-format include sub-intf** command or the **aaa intf-desc-format include adapter** command when the subscriber logs in.
- You can issue the **ipv6** keyword to display all IPv6 subscribers or include the IPv6 prefix to limit the display to only IPv6 subscribers on a specific network.
- You can issue the **icr-partition** keyword to display active subscribers for a particular ICR partition configured on a chassis.



NOTE: If you attempt to bring up tunneled subscribers on ACI-based VLAN subinterfaces on LAC devices with subscriber groups that are based on S-VLAN IDs (using the **ip vrrp vrid icr-partition group svlan** command on S-VLAN subinterfaces), the VLAN subinterface does not come up and a log message to denote its down state is not generated. If you attempt to bring up tunneled subscribers on ACI-based VLAN subinterfaces on LAC devices with subscriber groups that are based on VLAN IDs (using the **ip vrrp vrid icr-partition group vlan** command on VLAN subinterfaces), the subscribers over tunnels are brought up. However, on the LAC device, the subscribers are logged in outside of the ICR partition.

This behavior is expected when attempts are made to log in tunneled subscribers over ACI-based VLAN subinterfaces configured with ICR partitions with VLAN-based grouping or S-VLAN-based grouping.

- You can use the **profile** keyword to list subscribers who share the same profile.
- You can specify the **summary** keyword to display only summary information about active subscribers.

- In the Interface field in the output of the **show subscribers** command, for subscribers that are logged in to the router over VLAN interfaces configured on the LAG bundle using protocols such as DHCP or PPPoE, the logged-in subscriber name is displayed against the LAG bundle on the member interface where the user session is established. The subscriber sessions are displayed for the corresponding major interfaces, such as Ethernet, only if the subscribers are logged in over VLAN subinterfaces configured over major interfaces.

Action To display general subscriber information:

```
host1# show subscribers
```

```

Subscriber List
-----
User Name      Type      Addr|Endpt      Virtual
-----
fred           tst       10.10.65.86/radius default
bert           tst       192.168.10.3/user default
User Name      Interface
-----
fred           atm 2/1.42:100.104
bert           FastEthernet 5/2.4
User Name      Login Time      Circuit Id
-----
fred           06/05/12 10:58:42 atm 5/1.3
bert           06/05/12 10:59:08
User Name      Remote Id
-----
fred
bert           (800) 555-1212

```

To display detailed information about subscribers on the specified interface:

```
host1# show subscribers interface ethernet 5/2
```

```

Subscriber List
-----
User Name      Type      Addr|Endpt      Virtual
-----
bert           tst       192.168.10.3/user default
User Name      Interface
-----
bert           FastEthernet 5/2.4
User Name      Login Time      Circuit Id
-----
bert           06/05/12 10:59:08
User Name      Remote Id
-----
bert           (800) 555-0000

```

To display detailed information about subscribers on the specified LAG interface:

```
host1# show subscribers interface lag lag2.1:1-1
```

```

Subscriber List
-----
User Name      Type      Addr|Endpt      Router
-----

```

```

4101DHCPCLIENT@CT.NET      ip      2.0.0.3/user      default
User Name                    Interface
-----
4101DHCPCLIENT@CT.NET      lag lag2.1:1-1
User Name                    Login Time          Circuit Id
-----
4101DHCPCLIENT@CT.NET      09/10/29 02:07:51
User Name                    Remote Id
-----
4101DHCPCLIENT@CT.NET

```

To display detailed information about subscribers on the specified slot:

```

host1# show subscribers slot 5
Subscriber List
-----
User Name      Type      Addr|Endpt      Virtual
-----
fred           tst       10.10.65.86/radius default
User Name      Interface
-----
fred           atm 5/1.42:100.104
User Name      Login Time      Circuit Id
-----
fred           06/05/12 10:58:42 atm 5/1.3
User Name      Remote Id
-----
fred

```

To display detailed information about subscribers who share the same profile:

```

host1# show subscribers profile aaa
Subscriber List
-----
User Name      Type      Addr|Endpt      Virtual
-----
user           ppp       20.10.10.3/local default
user           ppp       20.10.10.8/local default
User Name      Interface
-----
user           FastEthernet 1/5
user           FastEthernet 1/5
User Name      Login Time      Circuit Id
-----
user           12/08/21 11:36:05
user           12/08/22 16:34:53
User Name      Remote Id      Profile Name
-----
user           aaa
user           aaa

```

To display the number of subscribers who share the same profile:

```

host1# show subscribers summary profile
Profile Name      Count
-----

```

```

aa                                2
aaa                              2
aab                              2
Total Subscribers : 6 (chassis-wide total)
Peak Subscribers : 6 (chassis-wide total)

```

To display the number of subscribers on each virtual router, as well as the total and peak subscribers for the chassis:

```

host1#show subscribers summary
Virtual
Router      Subscribers      Ppp      Ip      Tn1      Total
-----
default     1                1        0        0        1
Total Subscribers : 10 (chassis-wide total)
Peak Subscribers : 15 (chassis-wide total)

```

To display the number of subscribers on each port:

```

host1#show subscribers summary port
Interface      Count
-----
3/1            5
2/1            5
Total Subscribers : 10 (chassis-wide total)
Peak Subscribers : 15 (chassis-wide total)

```

To display the number of subscribers by domain name:

```

host1#show subscribers summary domain
Domain Name      Count
-----
abc.com          5
iii.com          5
Total Subscribers : 10 (chassis-wide total)
Peak Subscribers : 15 (chassis-wide total)

```

To display the number of subscribers by interface:

```

host1#show subscribers summary interface
Interface      Count
-----
ATM 3/2.1      1
ETHERNET 5/2.1 2
LAG lag1.100   1
Total Subscribers: 4 (chassis-wide total)
Peak Subscribers: 8 (chassis-wide total)

```

To display the number of subscribers by slot:

```

host1#show subscribers summary slot
Slot      Count
-----
3         1
5         4
Total Subscribers : 5 (chassis-wide total)
Peak Subscribers : 8 (chassis-wide total)

```

To display the number of subscribers by ICR partition:

```

host1#show subscribers summary icr-partition
ICR-Partition (location-id)      Count

```

```

-----
3/0.1.4                    5
3/0.2.5                    5
Total Subscribers: 10 (chassis-wide total)
Peak Subscribers: 15 (chassis-wide total)

```

To display the number of subscribers that are logged in on top of a LAG bundle:

```

host1#show subscribers summary lag
      Interface      Count
-----
LAG OLT              6
Total Subscribers : 6 (chassis-wide total)
Peak Subscribers : 6 (chassis-wide total)

```

Meaning [Table 35 on page 166](#) lists the **show subscribers** command output fields.

Table 35: show subscribers Output Fields

Field Name	Field Description
Username	Name of the subscriber NOTE: The complete username with the domain name (if available) is displayed regardless of the status of the strip domain feature on a virtual router or AAA domain map.
Type	Type of subscriber: atm, ip, ipsec, ppp, tnl (tunnel), or tst (test)
Addr Endpt	IP or IPv6 address and source of the address: l2tp, local, dhcp, radius, or user. For local, dhcp, radius, and user endpoints, the address is that of the user. When the endpoint is l2tp, the address is that of the LNS.
Virtual Router	Name of the virtual router context
Interface	Interface specifier over which the subscriber is connected
Login Time	Date, in YY/MM/DD format, and time the subscriber logged in
Circuit Id	User circuit ID value specified by PPPoE
Remote Id	User remote ID value specified by PPPoE
Total Subscribers	Number of active subscribers, chassis-wide
Peak Subscribers	Maximum value that is displayed in the Total Subscriber field during the time the router has been active, chassis-wide
Subscribers	Number of subscribers; the sum of the Ppp and Ip fields

Table 35: show subscribers Output Fields (*continued*)

Field Name	Field Description
Ppp	Number of PPPoA and PPPoE users, combined
Ip	Number of DHCP and IP subscriber manager users, combined
Tnl	Number of users tunneled to an LNS
Total	Total number of users per virtual router; the sum of the Ppp, Ip, and Tnl fields
Domain Name	Domain name used by the subscriber
ICR-Partition (location-id)	A unique identifier for each ICR partition on a chassis. Note that this ID is different from the partition name, which is configured using the ip vrrp vrid icr-partition partitionName command.
Count	Number of subscribers
Slot	Number of slot in the chassis

Related Documentation

- [show subscribers](#)

Monitoring Application Terminate Reason Mappings

Purpose Display information about the mappings for application terminate reasons.

Action To display the current terminate reasons that are mapped to a specific Acct-Terminate-Cause-Code:

This example uses the **radius** keyword to display all current terminate reasons mapped to RADIUS Acct-Terminate-Cause codes. The output lists all PPP mappings, followed by L2TP mappings, and then AAA mappings.

host1(config)#run show terminate-code radius

Apps	Terminate Reason	Description	Radius Code
-----	-----	-----	-----
ppp	authenticate-authenticator-timeout	authenticate authenticator timeout	17
ppp	authenticate-challenge-timeout	authenticate challenge timeout	10
ppp	authenticate-chap-no-resources	authenticate chap no resources	10
ppp	authenticate-chap-peer-authenticator-timeout	authenticate chap peer authenticator timeout	17
ppp	authenticate-deny-by-peer	authenticate deny by peer	17
ppp	authenticate-inactivity-timeout	authenticate inactivity timeout	4

```
ppp          authenticate-max-requests    authenticate max requests    10
--More--
```

To display all terminate reasons that are mapped to a specific terminate code:

This example uses the **radius** keyword and a RADIUS Acct-Terminate-Cause code (**radius 4**) to display all terminate reasons mapped to the specified terminate code.

```
host1(config)#run show terminate-code radius 4
```

Apps	Terminate Reason	Description	Radius Code
ppp	authenticate-inactivity-timeout	authenticate inactivity timeout	4
l2tp	session-timeout-inactivity	session timeout inactivity	4

To display all current mappings for a particular application's terminate reasons:

This example uses **aaa** as the application.

```
host1(config)#run show terminate-code aaa
```

Apps	Terminate Reason	Description	Radius Code
aaa	deny-server-not-available	deny server not available	17
aaa	deny-server-request-timeout	deny server request timed out	17
aaa	deny-authentication-failure	deny authentication failure from server	17
aaa	deny-address-assignment-failure	deny address assignment failure	17
aaa	deny-address-allocation-failure	deny address allocation failure	17
aaa	deny-no-address-allocation-resources	deny insufficient resources for address allocation	17
aaa	deny-unknown-subscriber	deny no such server entry	17
aaa	deny-no-resources	deny no resources available	10

--More--

To display the mapping for a specific terminate reason for an application:

This example uses **l2tp** as the application and **session-access-interface-down** as the terminate reason.

```
host1#show terminate-code l2tp session-access-interface-down
```

Terminate Reason	Description	Radius Code
session access interface down		8

Meaning Table 36 on page 168 lists the **show terminate-code** command output fields.

Table 36: show terminate-code Output Fields

Field Name	Field Description
Apps	The application generating the terminate reason; AAA, L2TP, PPP, or RADIUS client

Table 36: show terminate-code Output Fields (*continued*)

Field Name	Field Description
Terminate Reason	The application's terminate reason
Description	The terminate reason
Radius Code	The RADIUS Acct-Terminate-Cause code to which the application's terminate reason is mapped

Related Documentation • [show terminate-code](#)

Monitoring IPv6 Local Pools for DHCP Prefix Delegation By All Configured Pools

Purpose Display a summary of all the IPv6 local address pools configured on a virtual router, along with the prefix ranges in each of those pools, total number of prefixes that can be allocated to clients, and the number of prefixes that are in use by clients.

Action To display information about all the IPv6 local address pools configured on a virtual router:

host1#show ipv6 local pool

IPv6 Local Address Pools				
Pool	Start	End	Total	In Use
ipv6Pool-pppoa	2002:2002::/48	2002:2002:ffff::/48	65536	0
ipv6Pool-pppoe	3003:3003::/48	3003:3003:ffff::/48	65536	0
example	4004:4004::/48	4004:4004:ffff::/48	65536	16000

Meaning [Table 37 on page 169](#) lists the **show ipv6 local pool** command output fields.

Table 37: show ipv6 local pool Output Fields

Field Name	Field Description
Pool	Names of IPv6 local address pools configured on the virtual router
Start	Starting prefix of the range of prefixes configured in a particular pool
End	Ending prefix of the range of prefixes configured in a particular pool
Total	Number of prefixes available for allocation to clients from a particular pool
In Use	Number of prefixes in a pool that are currently used by DHCPv6 clients

Related Documentation • [show ipv6 local pool](#)

Monitoring IPv6 Local Pools for DHCP Prefix Delegation By Pool Name

Purpose Display prefix delegation details for an IPv6 local address pool configured on a virtual router.

Action To display prefix delegation information for a specific IPv6 local address pool:

host1#show ipv6 local pool example

Pool : example

Utilization : 24

Start	End	Total	In Use	Exclude	Util	Preferred Lifetime	Valid Lifetime
-----	-----	-----	-----	-----	-----	-----	-----
4004:4004::/48	4004:4004:ffff::/48	65536	16000	1	24	30 minutes	1 day
Exclude	4004:4004::/48						
Dns Servers	5:5:5:5:5:5:5:5						
	6:6:6:6:6:6:6:6						
Domain Search List	example-1.com						
	example-2.com						
	example-3.com						
	example-4.com						

Meaning [Table 38 on page 170](#) lists the **show ipv6 local pool poolName** command output fields.

Table 38: show ipv6 local pool poolName Output Fields

Field Name	Field Description
Pool	Name of the IPv6 local address pool for which prefix delegation details are displayed
Utilization	Percentage of IPv6 prefixes currently allocated to clients from the local address pool
Start	Starting prefix of the range of prefixes configured in a particular pool
End	Ending prefix of the range of prefixes configured in a particular pool
Total	Number of prefixes available for allocation to clients from a particular pool
In Use	Number of prefixes in a pool that are currently used by DHCPv6 clients
Preferred Lifetime	Amount of time for which the prefix remains preferred for the requesting router to use

Table 38: show ipv6 local pool poolName Output Fields (*continued*)

Field Name	Field Description
Valid Lifetime	Amount of time for which the prefix remains valid for the requesting router to use
Exclude	Prefix length or prefix range excluded from allocation to the requesting router
Util	Percentage of prefixes currently allocated to clients from a particular prefix range in the pool
Dns Servers	List of IPv6 addresses of DNS servers to be sent to clients in the DHCPv6 responses
Domain Search List	List of domain names configured in the IPv6 local pool for DNS resolution

Related Documentation

- [show ipv6 local pool](#)

Monitoring IPv6 Local Pool Statistics for DHCP Prefix Delegation

Purpose Display IPv6 local address pool statistics used for DHCP prefix delegation to requesting routers.

Action To display all IPv6 local address pool statistics for prefix delegation to clients:

```
host1#show ipv6 local pool statistics
IPv6 Local Address Pool Statistics
-----
Statistic      Value
-----
Allocations    0
Allocation Errors 0
Releases       0
Release Errors  0
```

Meaning [Table 39 on page 171](#) lists the **show ipv6 local pool statistics** command output fields.

Table 39: show ipv6 local pool statistics Output Fields

Field Name	Field Description
Allocations	Number of prefixes allocated to DHCPv6 clients from the local address pool
Allocation Errors	Number of errors encountered during the allocation of prefixes
Releases	Number of prefixes released back to the pool

Table 39: show ipv6 local pool statistics Output Fields (*continued*)

Field Name	Field Description
Release Errors	Number of errors encountered during the process of release of previously assigned prefixes by the requesting router

Related Documentation

- [show ipv6 local pool](#)

Monitoring IPv6 Local Pools for Neighbor Discovery Router Advertisements for all Configured Pools

Purpose Display a summary of all the IPv6 local address pools configured on a virtual router, along with the prefix ranges in each of those pools, total number of prefixes that can be allocated to clients, and the number of prefixes that are in use by clients for Neighbor Discovery router advertisements.

Action To display information about all the IPv6 local address pools configured on a virtual router for Neighbor Discovery router advertisements:

```
host1#show ipv6 local ndra-pool
```

IPv6 Local Address ND-RA Pools				
Pool	Start	End	Total	In Use
ipv6Pool-expm1	2002:2002::/64	2002:2002:ffff::/64	65536	0
ipv6Pool-expm2	3003:3003::/48	3003:3003:ffff::/48	65536	0
example	4004:4004:0:ff00::/64	4004:4004:ffff::/48	65536	16000

Meaning [Table 40 on page 172](#) lists the **show ipv6 local ndra-pool** command output fields.

Table 40: show ipv6 local ndra-pool Output Fields

Field Name	Field Description
Pool	Names of IPv6 Neighbor Discovery router advertisement local address pools configured on the virtual router
Start	Starting prefix of the range of prefixes configured in a particular Neighbor Discovery router advertisements pool
End	Ending prefix of the range of prefixes configured in a particular Neighbor Discovery router advertisements pool
Total	Number of prefixes available for allocation to clients from a particular Neighbor Discovery router advertisements pool

Table 40: show ipv6 local ndra-pool Output Fields (*continued*)

Field Name	Field Description
In Use	Number of prefixes in a pool that are currently used by Neighbor Discovery clients

Related Documentation

- [IPv6 Prefix Allocation Using Neighbor Discovery Router Advertisements from IPv6 Address Pools Overview on page 57](#)
- [Monitoring IPv6 Local Pools for Neighbor Discovery Router Advertisements by Pool Name on page 173](#)
- [Monitoring IPv6 Local Pool Statistics for Neighbor Discovery Router Advertisements Allocation of Prefixes on page 174](#)
- ***show ipv6 local ndra-pool***

Monitoring IPv6 Local Pools for Neighbor Discovery Router Advertisements by Pool Name

Purpose Display information about an IPv6 local address pool for Neighbor Discovery router advertisements configured on a virtual router.

Action To display information about an IPv6 local address pool for Neighbor Discovery router advertisements configured on a virtual router:

host1#show ipv6 local ndra-pool example

```
Pool : example
-----
Utilization : 24
```

Start	End	Total	In Use	Exclude	Util
-----	-----	-----	-----	-----	-----
2002:2002::/64	2002:2002:ffff::/64	65536	0	0	0
3003:3003::/64	3003:3003:0:1000::/64	17	0	0	0
4004:4004:0:ff00::/64	4004:4004:0:ffff::/64	256	0	0	0

Meaning [Table 41 on page 173](#) lists the **show ipv6 local ndra-pool *poolName*** command output fields.

Table 41: show ipv6 local ndra-pool poolName Output Fields

Field Name	Field Description
Pool	Names of IPv6 Neighbor Discovery router advertisements local address pools configured on the virtual router

Table 41: show ipv6 local ndra-pool poolName Output Fields (*continued*)

Field Name	Field Description
Start	Starting prefix of the range of prefixes configured in a particular Neighbor Discovery router advertisements pool
End	Ending prefix of the range of prefixes configured in a particular Neighbor Discovery router advertisements pool
Total	Number of prefixes available for allocation to clients from a particular Neighbor Discovery router advertisements pool
In Use	Number of prefixes in a pool that are currently used by Neighbor Discovery clients
Exclude	Prefix length or prefix range excluded
Util	Percentage of prefixes currently allocated to clients from a particular prefix range in the Neighbor Discovery router advertisements pool

Related Documentation

- [IPv6 Prefix Allocation Using Neighbor Discovery Router Advertisements from IPv6 Address Pools Overview on page 57](#)
- [Monitoring IPv6 Local Pools for Neighbor Discovery Router Advertisements for all Configured Pools on page 172](#)
- [Monitoring IPv6 Local Pool Statistics for Neighbor Discovery Router Advertisements Allocation of Prefixes on page 174](#)
- *show ipv6 local ndra-pool*

Monitoring IPv6 Local Pool Statistics for Neighbor Discovery Router Advertisements Allocation of Prefixes

Purpose Display IPv6 local address pool statistics used for Neighbor Discovery router advertisements to requesting routers.

Action To display all IPv6 local address pool statistics for Neighbor Discovery router advertisements to requesting routers:

```
host1#show ipv6 local ndra-pool statistics
IPv6 Local Address Pool Statistics
-----
Statistic      Value
-----
Allocations    0
Allocation Errors 0
Releases      0
Release Errors 0
```

Meaning [Table 42 on page 175](#) lists the **show ipv6 local ndra-pool statistics** command output fields.

Table 42: show ipv6 local ndra-pool statistics Output Fields

Field Name	Field Description
Allocations	Number of prefixes allocated to Neighbor Discovery router advertisements clients from the local address pool
Allocation Errors	Number of errors encountered during the allocation of prefixes
Releases	Number of prefixes released back to the pool
Release Errors	Number of errors encountered during the process of release of previously assigned prefixes by the requesting router

- Related Documentation**
- [IPv6 Prefix Allocation Using Neighbor Discovery Router Advertisements from IPv6 Address Pools Overview on page 57](#)
 - [Monitoring IPv6 Local Pools for Neighbor Discovery Router Advertisements for all Configured Pools on page 172](#)
 - [Monitoring IPv6 Local Pools for Neighbor Discovery Router Advertisements by Pool Name on page 173](#)
 - ***show ipv6 local ndra-pool***

Monitoring the Status of the Override Feature to Initiate IPCP and IPv6CP Negotiations Based on RADIUS-Returned Attributes

Purpose Displays the status of the override feature to initiate IPCP and IPv6CP negotiations for IPv4 and IPv6 clients based on RADIUS-returned attributes.

Action To display the status of the override feature to initiate IPCP and IPv6CP negotiations for IPv4 and IPv6 clients based on RADIUS-returned attributes:

```
host1#show aaa radius-override-ncp-negotiation
```

```
enabled
```

- Related Documentation**
- [Overriding AAA to Perform IPCP and IPv6CP Negotiations Based on RADIUS-Returned Attributes on page 107](#)
 - ***show aaa radius-override-ncp-negotiation***

PART 2

Managing RADIUS and TACACS+

- [Configuring RADIUS Attributes on page 179](#)
- [Configuring RADIUS Dynamic-Request Server on page 227](#)
- [Configuring RADIUS Relay Server on page 235](#)
- [RADIUS Attribute Descriptions on page 241](#)
- [Application Terminate Reasons on page 265](#)
- [Monitoring RADIUS on page 291](#)
- [Configuring TACACS+ on page 305](#)
- [Monitoring TACACS+ on page 315](#)

CHAPTER 4

Configuring RADIUS Attributes

This chapter identifies the Remote Authentication Dial-In User Service (RADIUS) attributes that JunosE Software supports and describes the RADIUS attributes you can configure with the command-line interface (CLI). RADIUS attributes are discussed in the following sections:

- [RADIUS Overview on page 180](#)
- [RADIUS Platform Considerations on page 181](#)
- [RADIUS References on page 181](#)
- [Subscriber AAA Access Messages Overview on page 182](#)
- [RADIUS IETF Attributes Supported for Subscriber AAA Access Messages on page 183](#)
- [Juniper Networks VSAs Supported for Subscriber AAA Access Messages on page 186](#)
- [Processing DNS Addresses from Microsoft RADIUS VSAs for PPP Clients During IPCP on page 191](#)
- [Subscriber AAA Accounting Messages Overview on page 192](#)
- [RADIUS IETF Attributes Supported for Subscriber AAA Accounting Messages on page 193](#)
- [Juniper Networks VSAs Supported for Subscriber AAA Accounting Messages on page 196](#)
- [RADIUS IETF Attributes Supported for AAA Tunnel Accounting Messages on page 200](#)
- [AAA Access Messages During IPCP Negotiations for Dual-Stack Subscribers on page 201](#)
- [AAA Accounting Messages During IPCP Negotiations for Dual-Stack Subscribers on page 203](#)
- [DSL Forum VSAs in AAA Access and Accounting Messages Overview on page 206](#)
- [DSL Forum VSAs Supported for AAA Access and Accounting Messages on page 206](#)
- [RADIUS Attributes Supported for CLI AAA Messages on page 208](#)
- [CLI Commands Used to Modify RADIUS Attributes on page 209](#)
- [CLI Commands Used to Configure RADIUS IETF Attributes on page 209](#)
- [CLI Commands Used to Configure Juniper Networks VSAs on page 213](#)
- [CLI Commands Used to Include ANCP-Related Juniper Networks VSAs in Access and Accounting Messages on page 215](#)
- [CLI Commands Used to Include DSL Forum VSAs in Access and Accounting Messages on page 217](#)

- [CLI Commands Used to Include or Exclude Attributes in RADIUS Messages on page 218](#)
- [CLI Commands Used to Ignore Attributes when Receiving Access-Accept Messages on page 223](#)
- [RADIUS Per-Profile Attribute List Configuration Overview on page 223](#)
- [Example: Configuring RADIUS-Specific Attributes on page 224](#)

RADIUS Overview

RADIUS is a distributed client/server that protects networks against unauthorized access. RADIUS clients running on a Juniper Networks E Series Broadband Services Router send authentication requests to a central RADIUS server.

You can access the RADIUS server through either a subscriber line or the CLI.



NOTE: For CLI/telnet users only—For CLI security, the router supports the RADIUS Access-Challenge message. The RADIUS server uses this message to send the user a challenge requiring a response. The router then displays the single reply message and attempts to authenticate the user with the new response as the password.

The central RADIUS server stores all the required user authentication and network access information. RADIUS informs the router of the privilege levels for which RADIUS-authenticated users have enable access. The router permits or denies enable access accordingly.

The RADIUS server is configured and managed by a RADIUS administrator. See your RADIUS server documentation for information about configuring and managing a RADIUS server.

The E Series RADIUS client uses the IP address in the router ID unless you explicitly set an IP address by using the *radius update-source-addr* command.

To explicitly set the source address, perform the following tasks:

- Configure the RADIUS update-source address.
- Set this address on the RADIUS server if required.



NOTE: For additional RADIUS information about topics such as restricting user access, vty line authentication, or SSH, see the *Passwords and Security* chapter in the *JunosE System Basics Configuration Guide*.

RADIUS Services

RADIUS provides three distinct services:

- Authentication—Determines whether or not a user is allowed to access a specific service or resource.
- Authorization—Associates connection attributes or characteristics with a specific user.
- Accounting—Tracks service use by subscribers.

RADIUS Attributes

JunosE Software supports the RADIUS attributes and vendor-specific attributes (VSAs) listed in this chapter. These attributes define specific authentication, authorization, and accounting elements in a user's profile. The profile is stored on the RADIUS server. RADIUS messages contain RADIUS attributes to communicate information between an E Series Broadband Services Router and the RADIUS server.

Note these guidelines about RADIUS attribute numbers:

- The number, such as [1], that appears in brackets before each attribute is the attribute's standard number.
- Any attribute number beginning with 26, such as [26-1], identifies a vendor-specific attribute.

Related Documentation

- [RADIUS Authentication and Accounting Servers Configuration Overview on page 17](#)
- [RADIUS Platform Considerations on page 181](#)
- [RADIUS IETF Attributes on page 241](#)

RADIUS Platform Considerations

RADIUS is supported on all E Series routers.

For information about the modules supported on E Series routers:

- See the *ERX Module Guide* for modules supported on ERX7xx models, ERX14xx models, and the ERX310 Broadband Services Router.
- See the *E120 and E320 Module Guide* for modules supported on the Juniper Networks E120 and E320 Broadband Services Routers.

Related Documentation

- [RADIUS Overview on page 180](#)

RADIUS References

For more information about RADIUS, consult the following resources:

- RFC 2865—Remote Authentication Dial In User Service (RADIUS) (June 2000)
- RFC 2866—RADIUS Accounting (June 2000)
- RFC 2867—RADIUS Accounting Modifications for Tunnel Protocol Support (June 2000)

- RFC 2868—RADIUS Attributes for Tunnel Protocol Support (June 2000)
- RFC 2869—RADIUS Extensions (June 2000)
- RFC 4679—DSL Forum Vendor-Specific RADIUS Attributes (September 2006)
- GSMP extensions for layer2 control (L2C) Topology Discovery and Line Configuration—draft-wadhwa-gsmp-l2control-configuration-00.txt (July 2006 expiration)

**Related
Documentation**

- [RADIUS Overview on page 180](#)
- [Subscriber AAA Access Messages Overview on page 182](#)
- [Subscriber AAA Accounting Messages Overview on page 192](#)
- [DSL Forum VSAs in AAA Access and Accounting Messages Overview on page 206](#)

Subscriber AAA Access Messages Overview

Authorization and authentication access messages identify subscribers before the RADIUS server grants or denies them access to the network or network services. When an application requests user authentication, the request must have certain authenticating attributes, such as a user's name, password, and the particular type of service the user is requesting. This information is sent in the authentication request via the RADIUS protocol to the RADIUS server. In response, the RADIUS server grants or denies the request.

The router supports the following types of authentication and authorization messages:

- Access-Request—Requests client authentication. RADIUS responds to a client authentication request with either an Access-Accept, an Access-Reject, or an Access-Challenge message. An Access-Request message can contain a number of RADIUS attributes.
- Access-Accept—Grants the client's access request and can provide specific configuration information necessary to begin delivery of service to the user.
- Access-Reject—Sent if any value of the received attributes is not acceptable.
- Access-Challenge—Sent to the client, requesting additional authentication information.
- Change-of-Authorization-Request (COA-Request)—Dynamically modifies session attributes, such as data filters.
- Disconnect-Request—Immediately terminates a user session.

**Related
Documentation**

- [RADIUS IETF Attributes Supported for Subscriber AAA Access Messages on page 183](#)
- [Juniper Networks VSAs Supported for Subscriber AAA Access Messages on page 186](#)

RADIUS IETF Attributes Supported for Subscriber AAA Access Messages

Table 43 on page 183 lists the Access-Request, Access-Accept, Access-Reject, Access-Challenge, COA, and Disconnect-Request attributes supported by JunosE Software. The following notes are referenced in Table 43 on page 183:

1. Attribute is used by Access-Request messages when terminating a PPP connection at the LNS or the initiating LAC.
2. Attribute is used to support pass-through exchange of EAP messages.
3. Attribute is used by Access-Challenge messages to set the PPP retransmission timeout used for EAP request packets.

Table 43 on page 183 lists the RADIUS IETF attributes supported for Access-Request, Access-Accept, Access-Reject, COA-Request, and Disconnect-Request messages.

Table 43: AAA Access Message RADIUS IETF Attributes Supported

Attribute Number	Attribute Name	Access-Request	Access-Accept	Access-Reject	Access-Challenge	COA-Request	Disconnect-Request
[1]	User-Name	✓	✓	–	–	✓	✓
[2]	User-Password	✓	–	–	–	–	–
[3]	CHAP-Password	✓	–	–	–	–	–
[4]	NAS-IP-Address	✓	–	–	–	–	–
[5]	NAS-Port	✓	–	–	–	–	–
[6]	Service-Type	✓	✓	–	–	–	–
[7]	Framed-Protocol	✓	✓	–	–	–	–
[8]	Framed-IP-Address	✓	✓	–	–	✓	–
[9]	Framed-IP-Netmask	–	✓	–	–	–	–
[11]	Filter-Id	–	✓	–	–	–	–
[12]	Framed-MTU (See Note 2.)	✓	✓	–	–	–	–
[18]	Reply-Message (See Note 2.)	–	✓	✓	✓	–	–
[22]	Framed-Route	–	✓	–	–	–	–
[24]	State (See Note 2.)	–	–	✓	✓	–	–

Table 43: AAA Access Message RADIUS IETF Attributes Supported
(continued)

Attribute Number	Attribute Name	Access-Request	Access-Accept	Access-Reject	Access-Challenge	COA-Request	Disconnect-Request
[25]	Class	–	✓	–	–	–	–
[27]	Session-Timeout (See Note 2.) (See Note 3.)	–	✓	✓	✓	–	–
[28]	Idle-Timeout	–	✓	–	–	–	–
[30]	Called-Station-Id	✓	–	–	–	–	–
[31]	Calling-Station-Id	✓	–	–	–	✓	–
[32]	NAS-Identifier	✓	–	–	–	–	–
[33]	Proxy-State	✓	–	–	–	–	–
[44]	Acct-Session-Id	✓	–	–	–	✓	✓
[50]	Acct-Multi-Session-Id	✓	–	–	–	–	✓
[60]	CHAP-Challenge	✓	–	–	–	–	–
[61]	NAS-Port-Type	✓	–	–	–	–	–
[62]	Port-Limit	–	✓	–	–	–	–
[64]	Tunnel-Type (See Note 1.)	✓	✓	–	–	–	–
[65]	Tunnel-Medium-Type (See Note 1.)	✓	✓	–	–	–	–
[66]	Tunnel-Client-Endpoint (See Note 1.)	✓	✓	–	–	–	–
[67]	Tunnel-Server-Endpoint (See Note 1.)	✓	✓	–	–	–	–
[68]	Acct-Tunnel-Connection (See Note 1.)	✓	–	–	–	–	–
[69]	Tunnel-Password	–	✓	–	–	–	–
[77]	Connect-Info	✓	–	–	–	–	–

Table 43: AAA Access Message RADIUS IETF Attributes Supported
(continued)

Attribute Number	Attribute Name	Access-Request	Access-Accept	Access-Reject	Access-Challenge	COA-Request	Disconnect-Request
[79]	EAP-Message (See Note 2.)	✓	✓	✓	✓	–	–
[80]	Message-Authenticator (See Note 2.)	✓	✓	✓	✓	–	–
[82]	Tunnel-Assignment-Id	–	✓	–	–	–	–
[83]	Tunnel-Preference	–	✓	–	–	–	–
[85]	Acct-Interim-Interval	–	✓	–	–	–	–
[87]	NAS-Port-Id	✓	–	–	–	✓	–
[88]	Framed-Pool	–	✓	–	–	–	–
[90]	Tunnel-Client-Auth-Id (See Note 1.)	✓	✓	–	–	–	–
[91]	Tunnel-Server-Auth-Id (See Note 1.)	✓	✓	–	–	–	–
[96]	Framed-Interface-Id	–	✓	–	–	–	–
[97]	Framed-Ipv6-Prefix	–	✓	–	–	–	–
[99]	Framed-Ipv6-Route	–	✓	–	–	–	–
[100]	Framed-IPv6-Pool	–	✓	–	–	–	–
[101]	Error-Cause	–	–	–	–	✓	✓
[123]	Delegated-IPv6-Prefix	–	✓	–	–	–	–
[135]	Ascend-Primary-Dns	–	✓	–	–	–	–
[136]	Ascend-Secondary-Dns	–	✓	–	–	–	–
[144]	DS-Lite-Tunnel-Name	–	✓	–	–	–	–
[188]	Ascend-Num-In-Multilink	✓	–	–	–	–	–
[242]	Ascend-Data-Filter	–	✓	–	–	–	–

- Related Documentation**
- [Subscriber AAA Access Messages Overview on page 182](#)
 - [CLI Commands Used to Configure RADIUS IETF Attributes on page 209](#)
 - [CLI Commands Used to Include or Exclude Attributes in RADIUS Messages on page 218](#)
 - [CLI Commands Used to Ignore Attributes when Receiving Access-Accept Messages on page 223](#)
 - [RADIUS IETF Attributes on page 241](#)

Juniper Networks VSAs Supported for Subscriber AAA Access Messages

Table 44 on page 186 lists the Juniper Networks (Vendor ID 4874) VSAs supported for Access-Request, Access-Accept, Access-Reject, COA-Request, and Disconnect-Request messages.

Table 44: AAA Access Message Juniper Networks (Vendor ID 4874) VSAs Supported

Attribute Number	Attribute Name	Access-Request	Access-Accept	Access-Reject	COA-Request	Disconnect-Request
[26-1]	Virtual-Router	–	✓	–	✓	–
[26-2]	Local-Address-Pool	–	✓	–	–	–
[26-3]	Local-Loopback-Interface	–	✓	–	–	–
[26-4]	Primary-DNS	–	✓	–	–	–
[26-5]	Secondary-DNS	–	✓	–	–	–
[26-6]	Primary-WINS (NBNS)	–	✓	–	–	–
[26-7]	Secondary-WINS (NBNS)	–	✓	–	–	–
[26-8]	Tunnel-Virtual-Router	–	✓	–	–	–
[26-9]	Tunnel-Password	–	✓	–	–	–
[26-10]	Ingress-Policy-Name	–	✓	–	–	–
[26-11]	Egress-Policy-Name	–	✓	–	–	–
[26-12]	Ingress-Statistics	–	✓	–	–	–
[26-13]	Egress-Statistics	–	✓	–	–	–
[26-14]	Service-Category	–	✓	–	–	–
[26-15]	PCR	–	✓	–	–	–

Table 44: AAA Access Message Juniper Networks (Vendor ID 4874) VSAs Supported (*continued*)

Attribute Number	Attribute Name	Access-Request	Access-Accept	Access-Reject	COA-Request	Disconnect-Request
[26-16]	SCR	–	✓	–	–	–
[26-17]	Mbs	–	✓	–	–	–
[26-22]	Sa-Validate	–	✓	–	–	–
[26-23]	IGMP-Enable	–	✓	–	–	–
[26-24]	Pppoe-Description	✓	–	–	–	–
[26-25]	Redirect-Vrouter-Name	–	✓	–	–	–
[26-26]	Qos-Profile-Name	–	✓	–	–	–
[26-30]	Tunnel-Nas-Port-Method	–	✓	–	–	–
[26-31]	SSC-Service-Bundle-Name	–	✓	–	–	–
[26-33]	Tunnel-Max-Sessions	–	✓	–	–	–
[26-34]	Framed-IP-Route-Tag	–	✓	–	–	–
[26-44]	Tunnel-Interface-ID	✓	–	–	–	–
[26-45]	Ipv6-Virtual-Router	–	✓	–	–	–
[26-46]	Ipv6-Local-Interface	–	✓	–	–	–
[26-47]	Ipv6-Primary-DNS	–	✓	–	–	–
[26-48]	Ipv6-Secondary-DNS	–	✓	–	–	–
[26-52]	RADIUS-Client-Address	✓	–	–	–	–
[26-53]	Service-Description	✓	–	–	–	–
[26-54]	L2tp-Recv-Window-Size	–	✓	–	–	–
[26-55]	DHCP-Options	✓	–	–	–	–
[26-56]	DHCP-MAC-Address	✓	–	–	–	–
[26-57]	DHCP-GI-Address	✓	–	–	–	–
[26-58]	LI-Action	–	✓	–	✓	–

Table 44: AAA Access Message Juniper Networks (Vendor ID 4874) VSAs Supported *(continued)*

Attribute Number	Attribute Name	Access-Request	Access-Accept	Access-Reject	COA-Request	Disconnect-Request
[26-59]	Med-Dev-Handle	–	✓	–	✓	–
[26-60]	Med-Ip-Address	–	✓	–	✓	–
[26-61]	Med-Port-Number	–	✓	–	✓	–
[26-62]	MLPPP-Bundle-Name	✓	–	–	–	–
[26-63]	Interface-Desc	✓	–	–	–	–
[26-64]	Tunnel-Group	–	✓	–	–	–
[26-65]	Activate-Service	–	✓	–	✓	–
[26-66]	Deactivate-Service	–	✓	–	✓	–
[26-67]	Service-Volume	–	✓	–	✓	–
[26-68]	Service-Timeout	–	✓	–	✓	–
[26-69]	Service-Statistics	–	✓	–	✓	–
[26-70]	Ignore-DF-Bit	–	✓	–	–	–
[26-71]	IGMP-Access-Name	–	✓	–	–	–
[26-72]	IGMP-Access-Src-Name	–	✓	–	–	–
[26-73]	IGMP-OIF-Map-Name	–	✓	–	–	–
[26-74]	MLD-Access-Name	–	✓	–	–	–
[26-75]	MLD-Access-Src-Name	–	✓	–	–	–
[26-76]	MLD-OIF-Map-Name	–	✓	–	–	–
[26-77]	MLD-Version	–	✓	–	–	–
[26-78]	IGMP-Version	–	✓	–	–	–
[26-79]	IP-Mcast-Adm-Bw-Limit	–	✓	–	–	–
[26-80]	IPv6-Mcast-Adm-Bw-Limit	–	✓	–	–	–
[26-81]	L2c-Information	✓	–	–	–	–

Table 44: AAA Access Message Juniper Networks (Vendor ID 4874) VSAs Supported *(continued)*

Attribute Number	Attribute Name	Access-Request	Access-Accept	Access-Reject	COA-Request	Disconnect-Request
[26-82]	QoS-Parameters	–	✓	–	–	–
[26-84]	Mobile-IP-Algorithm	–	✓	–	–	–
[26-85]	Mobile-IP-SPI	–	✓	–	–	–
[26-86]	Mobile-IP-Key	–	✓	–	–	–
[26-87]	Mobile-IP-Replay	–	✓	–	–	–
[26-88]	Mobile-IP-Access-Control-List	–	✓	–	–	–
[26-89]	Mobile-IP-Lifetime	–	✓	–	–	–
[26-90]	L2TP-Resynch-Method	–	✓	–	–	–
[26-91]	Tunnel-Switch-Profile	–	✓	–	–	–
[26-92]	L2C-Up-Stream-Data	✓	–	–	–	–
[26-93]	L2C-Down-Stream-Data	✓	–	–	–	–
[26-94]	Tunnel-Tx-Speed-Method	–	✓	–	–	–
[26-95]	IGMP-Query-Interval	–	✓	–	–	–
[26-96]	IGMP-Max-Resp-Time	–	✓	–	–	–
[26-97]	IGMP-Immediate-Leave	–	✓	–	–	–
[26-98]	MLD-Query-Interval	–	✓	–	–	–
[26-99]	MLD-Max-Resp-Time	–	✓	–	–	–
[26-100]	MLD-Immediate-Leave	–	✓	–	–	–
[26-106]	Ipv6-Ingress-Policy-Name	–	✓	–	–	–
[26-107]	Ipv6-Egress-Policy-Name	–	✓	–	–	–
[26-110]	Acc-Loop-Cir-Id	✓	–	–	–	–
[26-111]	Acc-Aggr-Cir-Id-Bin	✓	–	–	–	–
[26-112]	Acc-Aggr-Cir-Id-Asc	✓	–	–	–	–

Table 44: AAA Access Message Juniper Networks (Vendor ID 4874) VSAs Supported (*continued*)

Attribute Number	Attribute Name	Access-Request	Access-Accept	Access-Reject	COA-Request	Disconnect-Request
[26-113]	Act-Data-Rate-Up	✓	–	–	–	–
[26-114]	Act-Data-Rate-Dn	✓	–	–	–	–
[26-115]	Min-Data-Rate-Up	✓	–	–	–	–
[26-116]	Min-Data-Rate-Dn	✓	–	–	–	–
[26-117]	Att-Data-Rate-Up	✓	–	–	–	–
[26-118]	Att-Data-Rate-Dn	✓	–	–	–	–
[26-119]	Max-Data-Rate-Up	✓	–	–	–	–
[26-120]	Max-Data-Rate-Dn	✓	–	–	–	–
[26-121]	Min-LP-Data-Rate-Up	✓	–	–	–	–
[26-122]	Min-LP-Data-Rate-Dn	✓	–	–	–	–
[26-123]	Max-Interlv-Delay-Up	✓	–	–	–	–
[26-124]	Act-Interlv-Delay-Up	✓	–	–	–	–
[26-125]	Max-Interlv-Delay-Dn	✓	–	–	–	–
[26-126]	Act-Interlv-Delay-Dn	✓	–	–	–	–
[26-127]	DSL-Line-State	✓	–	–	–	–
[26-128]	DSL-Type	✓	–	–	–	–
[26-129]	Ipv6-NdRa-Prefix	–	✓	–	–	–
[26-130]	QoS-Interfaceset-Name	–	✓	–	–	–
[26-140]	Service-Interim-Acct-Interval	–	✓	–	✓	–
[26-141]	Downstream-Calculated-Qos-Rate	✓	✓	–	✓	–
[26-142]	Upstream-Calculated-Qos-Rate	✓	✓	–	✓	–
[26-143]	Max-Clients-Per-Interface	–	✓	–	–	–
[26-144]	PPP-Monitor-Ingress-Only	–	✓	–	–	–

Table 44: AAA Access Message Juniper Networks (Vendor ID 4874) VSAs Supported (*continued*)

Attribute Number	Attribute Name	Access-Request	Access-Accept	Access-Reject	COA-Request	Disconnect-Request
[26-147]	Backup-Address-Pool	—	✓	—	—	—
[26-150]	ICR-Partition-Id	✓	—	—	—	—
[26-157]	Ipv6-Ndra-Pool	—	✓	—	—	—
[26-159]	DHCP-Option 82	✓	—	—	✓	—
[26-161]	Delegated-Ipv6-Pool	—	✓	—	—	—
[26-164]	Ipv4-release-control	✓	—	—	—	—
[26-165]	PCP-Server-Name	—	✓	—	—	—

Related Documentation

- [Subscriber AAA Access Messages Overview on page 182](#)
- [CLI Commands Used to Configure Juniper Networks VSAs on page 213](#)
- [CLI Commands Used to Include ANCP-Related Juniper Networks VSAs in Access and Accounting Messages on page 215](#)
- [CLI Commands Used to Include or Exclude Attributes in RADIUS Messages on page 218](#)
- [CLI Commands Used to Ignore Attributes when Receiving Access-Accept Messages on page 223](#)
- [Juniper Networks VSAs on page 248](#)

Processing DNS Addresses from Microsoft RADIUS VSAs for PPP Clients During IPCP

The RADIUS client, which is a B-RAS router, supports the processing and parsing of Microsoft RADIUS VSAs for the primary and secondary DNS addresses that are returned in the Access-Accept messages from the RADIUS server in an environment that contains PPP clients. The AAA application running on the router, which is the RADIUS client, transmits the DNS addresses to the PPP application in the authentication response message. PPP includes these DNS addresses in the Internet Protocol Control Protocol (IPCP) packets that are negotiated between the PPP client and the router.

The RADIUS client services the Microsoft vendor ID, 311, and does not discard the DNS server addresses that the Microsoft VSAs contain in the Access-Accept messages.

The PPP application uses Link Control Protocol (LCP) negotiations to establish the connection with the subscriber. PPP sends Network Control Protocol (NCP) packets to establish and configure the session with the client. After a link has been established and optional facilities have been negotiated as needed by the Link Control Protocol (LCP) between the customer premises equipment (CPE) and the provider edge (PE) device, PPP running on the PE device or the B-RAS server sends Network Control Protocol (NCP)

packets. When the CPE sends an IPCP negotiation, it negotiates IPv4 addresses, IPv6 addresses, or both.

After the PE device or the router receives an IPCP configuration request from the CPE, which starts the IPCP negotiation process, the B-RAS application running on the router requests a new IPv4 address from the RADIUS server. After successful authentication, the RADIUS server sends the Access-Accept message with all of the supported attributes for all established sessions.

If the Access-Accept message contains the MS-Primary-DNS-Server [311-28] and MS-Secondary-DNS-Server [311-29] RADIUS VSA attributes, which denote the primary and secondary DNS server addresses that can be requested by PPP clients from the B-RAS server during IPCP negotiations, the RADIUS client or the B-RAS server sends the values of the VSAs to the CPE in the IPCP packet that is negotiated.

During IPCPv4 negotiations, if the Access-Accept message contains both the Juniper Networks VSAs related to the DNS addresses (Primary-DNS [26-4] and Secondary-DNS [26-5]) and the Microsoft VSAs related to DNS addresses (MS-Primary-DNS-Server [311-28] and the MS-Secondary-DNS-Server [311-29]), the Juniper Networks VSAs take precedence over the Microsoft VSAs.

During IPCPv6 negotiations, if the Access-Accept message contains both the Juniper Networks VSAs related to the DNS addresses (Ipv6-Primary-DNS [26-47] and Ipv6-Secondary-DNS [26-48]) and the Microsoft VSAs related to DNS addresses (MS-Primary-DNS-Server [311-28] and the MS-Secondary-DNS-Server [311-29]), the Juniper Networks VSAs take precedence over the Microsoft VSAs.

With the capability to validate the Microsoft VSAs for primary and secondary DNS addresses enabled, the order of precedence of the RADIUS attributes in the Access-Accept messages to be used for IPCP negotiations is as follows:

1. Juniper Networks VSAs
2. RADIUS IETF attributes
3. Microsoft VSAs

Related Documentation

- *Processing NCP Negotiations in a Dual-Stack Environment Overview*

Subscriber AAA Accounting Messages Overview

Accounting messages identify service provisions and use on a per-user or per-tunnel basis. These messages keep track of when a particular service is initiated and terminated for a specific user.

JunosE Software supports the Acct-On message on startup or configuration of the first accounting server. Acct-Off messages are supported when the last RADIUS accounting server in a virtual router is removed, when the router is shut down, and when a virtual router that has configured RADIUS accounting servers is deleted.

Beginning with JunosE Release 11.0.0, you can configure the router to send the Partition-Accounting-On and Partition-Accounting-Off messages to the RADIUS server whenever an ICR partition toggles between the backup and master states.

The router supports the following types of accounting messages:

- Acct-Start
- Acct-Stop
- Interim-Acct
- Acct-On
- Acct-Off
- Partition-Accounting-On
- Partition-Accounting-Off

**Related
Documentation**

- [RADIUS IETF Attributes Supported for Subscriber AAA Accounting Messages on page 193](#)
- [Juniper Networks VSAs Supported for Subscriber AAA Accounting Messages on page 196](#)

RADIUS IETF Attributes Supported for Subscriber AAA Accounting Messages

[Table 45 on page 194](#) lists the RADIUS IETF attributes supported for Acct-Start, Acct-Stop, Interim-Acct, Acct-On, and Acct-Off messages.

The following notes are referred to in [Table 45 on page 194](#):

1. The attribute is used when terminating a PPP connection at the LNS or the initiating LAC.
2. For this attribute to be included, an IP address must be assigned to the subscriber.
3. The attribute is not included in Acct-Stop messages that are sent when a user session does not get established in one of the following situations.
 - The **aaa accounting acct-stop on-access-deny** command is enabled and the authentication server sends an Access-Reject (deny) message.
 - The **aaa accounting acct-stop on-aaa-failure** command is enabled and the authentication server issues an Access-Accept message (grant), but the AAA configuration denies access for the user. The **aaa accounting acct-stop on-aaa-failure** is enabled by default.
 - The **aaa accounting acct-stop on-aaa-failure** command is enabled and the user terminates before AAA receives the authentication response from the authentication server.
4. For this attribute to be included, an IPv6 interface ID must be assigned to the subscriber.
5. For this attribute to be included, at least one IPv6 prefix must be assigned to the subscriber.

Table 45: AAA Accounting Message RADIUS IETF Attributes Supported

Attribute Number	Attribute Name	Acct-Start	Acct-Stop	Interim-Acct	Acct-On	Acct-Off
[1]	User-Name	✓	✓	✓	–	–
[4]	NAS-IP-Address	✓	✓	✓	✓	✓
[5]	NAS-Port	✓	✓	✓	–	–
[6]	Service-Type	✓	✓	✓	–	–
[7]	Framed-Protocol (See Note 3.)	✓	✓	✓	–	–
[8]	Framed-IP-Address (See Note 2.)	✓	✓	✓	–	–
[9]	Framed-IP-Netmask	✓	✓	✓	–	–
[13]	Framed-Compression (See Note 3.)	✓	✓	✓	–	–
[22]	Framed-Route	✓	✓	✓	–	–
[25]	Class	✓	✓	✓	–	–
[30]	Called-Station-Id	✓	✓	✓	–	–
[31]	Calling-Station-Id	✓	✓	✓	–	–
[32]	NAS-Identifier	✓	✓	✓	✓	✓
[40]	Acct-Status-Type	✓	✓	✓	✓	✓
[41]	Acct-Delay-Time	✓	✓	✓	✓	✓
[42]	Acct-Input-Octets	–	✓	✓	–	–
[43]	Acct-Output-Octets	–	✓	✓	–	–
[44]	Acct-Session-Id	✓	✓	✓	✓	✓
[45]	Acct-Authentic	✓	✓	✓	✓	✓
[46]	Acct-Session-Time	–	✓	✓	–	–
[47]	Acct-Input-Packets	–	✓	✓	–	–
[48]	Acct-Output-Packets	–	✓	✓	–	–

Table 45: AAA Accounting Message RADIUS IETF Attributes Supported (*continued*)

Attribute Number	Attribute Name	Acct-Start	Acct-Stop	Interim-Acct	Acct-On	Acct-Off
[49]	Acct-Terminate-Cause	–	✓	–	–	✓
[50]	Acct-Multi-Session-Id (See Note 3.)	✓	✓	✓	–	–
[51]	Acct-Link-Count (See Note 3.)	✓	✓	✓	–	–
[52]	Acct-Input-Gigawords	–	✓	✓	–	–
[53]	Acct-Output-Gigawords	–	✓	✓	–	–
[55]	Event-Timestamp	✓	✓	✓	✓	✓
[61]	NAS-Port-Type	✓	✓	✓	–	–
[64]	Tunnel-Type (See Note 1.)	✓	✓	✓	–	–
[65]	Tunnel-Medium-Type (See Note 1.)	✓	✓	✓	–	–
[66]	Tunnel-Client-Endpoint (See Note 1.)	✓	✓	✓	–	–
[67]	Tunnel-Server-Endpoint (See Note 1.)	✓	✓	✓	–	–
[68]	Acct-Tunnel-Connection (See Note 1.)	✓	✓	✓	–	–
[77]	Connect-Info	✓	✓	✓	–	–
[82]	Tunnel-Assignment-Id (LAC only) (See Note 1.)	✓	✓	✓	–	–
[83]	Tunnel-Preference (LAC only)	✓	✓	✓	–	–
[87]	NAS-Port-Id	✓	✓	✓	–	–
[90]	Tunnel-Client-Auth-Id (See Note 1.)	✓	✓	✓	–	–
[91]	Tunnel-Server-Auth-Id (See Note 1.)	✓	✓	✓	–	–
[96]	Framed-Interface-Id (See Note 1.)	✓	✓	✓	–	–

Table 45: AAA Accounting Message RADIUS IETF Attributes Supported (*continued*)

Attribute Number	Attribute Name	Acct-Start	Acct-Stop	Interim-Acct	Acct-On	Acct-Off
[97]	Framed-Ipv6-Prefix (See Note 5.)	✓	✓	✓	–	–
[99]	Framed-IPv6-Route	✓	✓	✓	–	–
[100]	Framed-IPv6-Pool	✓	✓	✓	–	–
[123]	Delegated-Ipv6-Prefix	✓	✓	✓	–	–
[144]	DS-Lite-Tunnel-Name	✓	✓	✓	–	–
[188]	Ascend-Num-In-Multilink (See Note 3.)	✓	✓	✓	–	–

**Related
Documentation**

- [Subscriber AAA Accounting Messages Overview on page 192](#)
- [CLI Commands Used to Configure RADIUS IETF Attributes on page 209](#)
- [CLI Commands Used to Include or Exclude Attributes in RADIUS Messages on page 218](#)
- [RADIUS IETF Attributes on page 241](#)

Juniper Networks VSAs Supported for Subscriber AAA Accounting Messages

Table 46 on page 197 lists the Juniper Networks (Vendor ID 4874) VSAs supported for Acct-Start, Acct-Stop, Interim-Acct, Acct-On, Acct-Off, Partition-Accounting-On, and Partition-Accounting-Off messages.

The following notes are referred to in Table 46 on page 197:

1. The attribute is not included in Acct-Stop messages that are sent when a user session does not get established in one of the following situations.
 - The **aaa accounting acct-stop on-access-deny** command is enabled and the authentication server sends an Access-Reject (deny) message.
 - The **aaa accounting acct-stop on-aaa-failure** command is enabled and the authentication server issues an Access-Accept message (grant), but the AAA configuration denies access for the user. The **aaa accounting acct-stop on-aaa-failure** is enabled by default.
 - The **aaa accounting acct-stop on-aaa-failure** command is enabled and the user terminates before AAA receives the authentication response from the authentication server.
2. ERX routers send IPv6 accounting attributes in the Acct-Stop and Interim-Acct messages (stop, interim) when they are configured to return these attributes and

when the subscriber is either an IPv6 subscriber or a combined IPv4/IPv6 subscriber in a dual stack. For an IPv4 subscriber, IPv6 accounting attributes are not included in the accounting messages even if the IPv6 accounting is enabled.

In JunosE Release 10.1.x and lower-numbered releases, the combined accounting statistics were retrieved at the layer 2. Therefore, error or discarded packets in the layer 2 itself were excluded in these statistics. Because the layer 2 cannot detect the error or discarded packets in the layer 3, the combined statistics also include the error or discarded packets of the layer 3. In this release, with the support for RADIUS VSAs for IPv6 accounting, the IPv6 statistics are retrieved at the layer 3. To be consistent with the combined statistics, the error or discarded packets of the layer 3 are also included in these IPv6 statistics.

3. The ICR partition accounting messages comprise the following:

- Partition-Accounting-On—Sent to the RADIUS server whenever an ICR partition changes to the master state from the backup state. The Partition-Accounting-On message has the same Acct-Status-Type attribute value as the Accounting-On message, but also contains the ICR-Partition-Id VSA, which specifies the ICR partition to which this message corresponds.
- Partition-Accounting-Off—Sent to the RADIUS server when the partition changes from the master state to the backup state. However, in the event of a complete chassis failure, the Partition-Accounting-Off message is not sent. Partition-Accounting-Off message has the same Acct-Status-Type attribute value as the Accounting-Off message and contains the ICR-Partition-Id VSA to denote the ICR partition to which the message is associated.

For more information about how to configure and use ICR partitions, see the *Managing Interchassis Redundancy* chapter in the *JunosE Services Availability Configuration Guide*.

Table 46: AAA Accounting Message Juniper Network (Vendor ID 4874) VSAs Supported

Attribute Number	Attribute Name	Acct-Start	Acct-Stop	Interim-Acct	Acct-On	Acct-Off	Partition-Accounting-On	Partition-Accounting-Off
[26-10]	Ingress-Policy-Name	✓	✓	✓	–	–	–	–
[26-11]	Egress-Policy-Name	✓	✓	✓	–	–	–	–
[26-24]	Pppoe-Description (See Note 1.)	✓	✓	✓	–	–	–	–
[26-26]	QoS-Profile-Name	✓	✓	✓	–	–	–	–
[26-42]	Acct-Input-Gigapackets	–	✓	✓	–	–	–	–
[26-43]	Acct-Output-Gigapackets	–	✓	✓	–	–	–	–
[26-44]	Tunnel-Interface-Id	✓	✓	✓	–	–	–	–
[26-45]	Ipv6-Virtual-Router	✓	✓	✓	–	–	–	–

Table 46: AAA Accounting Message Juniper Network (Vendor ID 4874) VSAs Supported (continued)

Attribute Number	Attribute Name	Acct-Start	Acct-Stop	Interim-Acct	Acct-On	Acct-Off	Partition-AccountingOn	Partition-AccountingOff
[26-46]	Ipv6-Local-Interface	✓	✓	✓	–	–	–	–
[26-47]	Ipv6-Primary-DNS	✓	✓	✓	–	–	–	–
[26-48]	Ipv6-Secondary-DNS	✓	✓	✓	–	–	–	–
[26-51]	Disconnect-Cause	–	✓	–	–	–	–	–
[26-53]	Service-Description	✓	✓	✓	–	–	–	–
[26-55]	DHCP-Options (See Note 1.)	✓	✓	✓	–	–	–	–
[26-56]	DHCP-MAC-Address (See Note 1.)	✓	✓	✓	–	–	–	–
[26-57]	DHCP-GI-Address (See Note 1.)	✓	✓	✓	–	–	–	–
[26-62]	MLPPP-Bundle-Name	✓	✓	✓	–	–	–	–
[26-63]	Interface-Description	✓	✓	✓	–	–	–	–
[26-92]	L2C-Up-Stream-Data	✓	✓	✓	–	–	–	–
[26-93]	L2C-Down-Stream-Data	✓	✓	✓	–	–	–	–
[26-106]	Ipv6-Ingress-Policy-Name	✓	✓	✓	–	–	–	–
[26-107]	Ipv6-Egress-Policy-Name	✓	✓	✓	–	–	–	–
[26-110]	Acc-Loop-Cir-Id	✓	✓	✓	–	–	–	–
[26-111]	Acc-Aggr-Cir-Id-Bin	✓	✓	✓	–	–	–	–
[26-112]	Acc-Aggr-Cir-Id-Asc	✓	✓	✓	–	–	–	–
[26-113]	Act-Data-Rate-Up	✓	✓	✓	–	–	–	–
[26-114]	Act-Data-Rate-Dn	✓	✓	✓	–	–	–	–
[26-115]	Min-Data-Rate-Up	✓	✓	✓	–	–	–	–
[26-116]	Min-Data-Rate-Dn	✓	✓	✓	–	–	–	–
[26-117]	Att-Data-Rate-Up	✓	✓	✓	–	–	–	–

Table 46: AAA Accounting Message Juniper Network (Vendor ID 4874) VSAs Supported (continued)

Attribute Number	Attribute Name	Acct-Start	Acct-Stop	Interim-Acct	Acct-On	Acct-Off	Partition-AccountingOn	Partition-AccountingOff
[26-118]	Att-Data-Rate-Dn	✓	✓	✓	–	–	–	–
[26-119]	Max-Data-Rate-Up	✓	✓	✓	–	–	–	–
[26-120]	Max-Data-Rate-Dn	✓	✓	✓	–	–	–	–
[26-121]	Min-LP-Data-Rate-Up	✓	✓	✓	–	–	–	–
[26-122]	Min-LP-Data-Rate-Dn	✓	✓	✓	–	–	–	–
[26-123]	Max-Interlv-Delay-Up	✓	✓	✓	–	–	–	–
[26-124]	Act-Interlv-Delay-Up	✓	✓	✓	–	–	–	–
[26-125]	Max-Interlv-Delay-Dn	✓	✓	✓	–	–	–	–
[26-126]	Act-Interlv-Delay-Dn	✓	✓	✓	–	–	–	–
[26-127]	DSL-Line-State	✓	✓	✓	–	–	–	–
[26-128]	DSL-Type	✓	✓	✓	–	–	–	–
[26-129]	Ipv6-NdRa-Prefix	✓	✓	✓	–	–	–	–
[26-150]	ICR-Partition-Id (See Note 3.)	✓	✓	✓	–	–	✓	✓
[26-151]	Ipv6-Acct-Input-Octets (See Note 2.)	–	✓	✓	–	–	–	–
[26-152]	Ipv6-Acct-Output-Octets (See Note 2.)	–	✓	✓	–	–	–	–
[26-153]	Ipv6-Acct-Input-Packets (See Note 2.)	–	✓	✓	–	–	–	–
[26-154]	Ipv6-Acct-Output-Packets (See Note 2.)	–	✓	✓	–	–	–	–
[26-155]	Ipv6-Acct-Input-Gigawords (See Note 2.)	–	✓	✓	–	–	–	–
[26-156]	Ipv6-Acct-Output-Gigawords (See Note 2.)	–	✓	✓	–	–	–	–
[26-159]	DHCP-Option 82 (See Note 1.)	✓	✓	✓	–	–	–	–

Table 46: AAA Accounting Message Juniper Network (Vendor ID 4874) VSAs Supported (continued)

Attribute Number	Attribute Name	Acct-Start	Acct-Stop	Interim-Acct	Acct-On	Acct-Off	Partition-AccountingOn	Partition-AccountingOff
[26-164]	Ipv4-release-control	–	–	✓	–	–	–	–
[26-165]	PCP-Server-Name	✓	✓	✓	–	–	–	–

Related Documentation

- [Subscriber AAA Accounting Messages Overview on page 192](#)
- [CLI Commands Used to Configure Juniper Networks VSAs on page 213](#)
- [CLI Commands Used to Include ANCP-Related Juniper Networks VSAs in Access and Accounting Messages on page 215](#)
- [CLI Commands Used to Include or Exclude Attributes in RADIUS Messages on page 218](#)
- [Juniper Networks VSAs on page 248](#)

RADIUS IETF Attributes Supported for AAA Tunnel Accounting Messages

Table 47 on page 200 lists RADIUS attributes supported by the following tunnel-related accounting messages:

- Acct-Tunnel-Start
- Acct-Tunnel-Stop
- Acct-Tunnel-Reject
- Acct-Tunnel-Link-Start
- Acct-Tunnel-Link-Stop
- Acct-Tunnel-Link-Reject

Table 47: AAA Accounting Tunnel Message RADIUS Attributes Supported

Attribute Number	Attribute Name	Acct-Tunnel-Start	Acct-Tunnel-Stop	Acct-Tunnel-Reject	Acct-Tunnel-Link-Start	Acct-Tunnel-Link-Stop	Acct-Tunnel-Link-Reject
[1]	User-Name	–	–	–	✓	✓	–
[4]	NAS-IP-Address	✓	✓	✓	✓	✓	✓
[26-51]	Disconnect-Cause	–	–	–	–	✓	–
[32]	NAS-Identifier	✓	✓	✓	✓	✓	✓
[40]	Acct-Status-Type	✓	✓	✓	✓	✓	✓
[41]	Acct-Delay-Time	✓	✓	✓	✓	✓	✓

Table 47: AAA Accounting Tunnel Message RADIUS Attributes Supported *(continued)*

Attribute Number	Attribute Name	Acct-Tunnel-Start	Acct-Tunnel-Stop	Acct-Tunnel-Reject	Acct-Tunnel-Link-Start	Acct-Tunnel-Link-Stop	Acct-Tunnel-Link-Reject
[44]	Acct-Session-Id	✓	✓	✓	✓	✓	✓
[46]	Acct-Session-Time	–	✓	–	–	✓	–
[49]	Acct-Terminate-Cause	–	✓	✓	–	✓	✓
[55]	Event-Timestamp	✓	✓	✓	✓	✓	✓
[64]	Tunnel-Type	✓	✓	✓	✓	✓	✓
[65]	Tunnel-Medium-Type	✓	✓	✓	✓	✓	✓
[66]	Tunnel-Client-Endpoint	✓	✓	✓	✓	✓	✓
[67]	Tunnel-Server-Endpoint	✓	✓	✓	✓	✓	✓
[68]	Acct-Tunnel-Connection	✓	✓	✓	✓	✓	✓
[82]	Tunnel-Assignment-Id (LAC only)	✓	✓	✓	✓	✓	✓
[83]	Tunnel-Preference (LAC only)	–	–	–	✓	✓	✓
[86]	Acct-Tunnel-Packets-Lost	–	–	–	–	✓	✓
[90]	Tunnel-Client-Auth-Id	✓	✓	✓	✓	✓	✓
[91]	Tunnel-Server-Auth-Id	✓	✓	✓	✓	✓	✓

Related Documentation

- [Subscriber AAA Accounting Messages Overview on page 192](#)
- [RADIUS IETF Attributes on page 241](#)
- [Juniper Networks VSAs on page 248](#)

AAA Access Messages During IPCP Negotiations for Dual-Stack Subscribers

The `ipv4-release-control` RADIUS VSA attribute [26-164] can be configured to be sent in the Access-Request and Interim-Acct messages. You can use the `aaa ipv4 addr-saving` command to configure this attribute to be sent in the access and accounting messages and enable the PPP application to inform the RADIUS server about the released IPv4

address for dual-stack subscribers, immediately after the address is released. The following sections describe the different scenarios during the negotiation of IPCP packets for IPv4 addresses, and the transmission of access messages between the router and the AAA server.

Access-Request Messages When an IPv4 Address is Renegotiated

During IPCP renegotiation of IPv4 addresses, the router always includes the following RADIUS attributes in the Access-Request messages sent to the RADIUS server:

- [1] User-Name
- [2] User-Password
- [4] NAS-Ip-Address
- [5] NAS-Port
- [25] Class
- [44] Acct-Session-Id
- [87] NAS-Port-Id
- [32] NAS-Identifier
- [26-164] Ipv4-release-control

The other attributes that are supported for this message are optional during renegotiations.

Access-Accept Messages When an IPv4 Address is Assigned

When the Access-Accept message is sent from the RADIUS server after the subscriber is successfully authenticated, during the initial session establishment, this message might or might not contain the Framed-Ip-Address [8] attribute.

When an IPv4 address is delegated to the CPE during an IPCP negotiation, the following RADIUS attributes are always included in the Access-Accept messages:

- [8] Framed-Ip-Address
- [9] Framed-Ip-Netmask
- [25] Class

The following RADIUS attributes can be optionally included in the Access-Accept messages based on the user topology settings:

- [242] Ascend-Data-Filter (always included if the policy is defined using this attribute for clients)
- [26-4] Primary-DNS
- [26-5] Secondary-DNS (included if the B-RAS user's DNS server is different from the previously stored entry during IPCP negotiation)

The following RADIUS attributes are never included in the Access-Accept messages when this functionality to optimally utilize IPv4 addresses is configured:

- [26-65] Activate-Service
- [242] Ascend-Data-Filter
- [26-66] Deactivate-Service
- [26-58] LI-Action
- [27] Session-Timeout
- [28] Idle-Timeout
- [97] Framed-Ipv6-Prefix
- [123] Delegated-Ipv6-Prefix

**Related
Documentation**

- [AAA Accounting Messages During IPCP Negotiations for Dual-Stack Subscribers on page 203](#)
- [Juniper Networks VSAs on page 248](#)
- [Juniper Networks VSAs Supported for Subscriber AAA Access Messages on page 186](#)
- [*IPCP Renegotiation of IPv4 Addresses for Dual-Stack Subscribers*](#)

AAA Accounting Messages During IPCP Negotiations for Dual-Stack Subscribers

The Ipv4-release-control RADIUS VSA attribute [26-164] is supported only in Interim-Acct messages. The following sections describe the different scenarios during the negotiation of IPCP packets for IPv4 addresses and release of addresses for terminated user sessions in a dual-stack network, and the transmission of interim accounting messages between the router and the AAA server.

Interim-Acct Messages When an IPv4 Address is Assigned

The following RADIUS attributes are always included in Interim-Acct messages when an IPv4 address is assigned to a dual-stack subscriber during an IPCP negotiation:

- [55] Event-Timestamp
- [25] Class
- [97] Framed-Ipv6-Prefix
- [8] Framed-Ip-Address
- [9] Framed-Ip-Netmask

The other attributes supported in Interim-Acct messages are included, based on the user configuration settings.

Interim-Acct Messages When an IPv4 Address is Not Assigned

The following RADIUS attributes are always included in Interim-Acct messages when an IPv4 address is not assigned to a dual-stack subscriber during an IPCP negotiation:

- [55] Event-Timestamp
- [25] Class
- [97] Framed-Ipv6-Prefix

The other attributes supported in Interim-Acct messages are included, based on the user configuration settings.

The following RADIUS attributes are never included in Interim-Acct messages when IPv4 address assignment fails:

- [8] Framed-Ip-Address
- [9] Framed-Ip-Netmask

Interim-Acct Messages During a Successful IPCP Renegotiation for IPv4 Addresses

The following RADIUS attributes are always included in Interim-Acct messages upon the receipt of IPCP configuration requests for IPv4 addresses from CPE and a successful IPCP renegotiation:

- [26-164] Ipv4-release-control
- [55] Event-Timestamp
- [25] Class
- [97] Framed-Ipv6-Prefix
- [8] Framed-Ip-Address
- [9] Framed-Ip-Netmask
- [123] Delegated-Ipv6-Prefix
- [44] Acct-Session-Id

The other attributes supported in Interim-Acct messages are included, based on the user configuration settings.

Interim-Acct Messages During a Failed IPCP Renegotiation for IPv4 Addresses

The following RADIUS attributes are always included in Interim-Acct messages after the receipt of IPCP configuration requests for IPv4 addresses from CPE and a failed IPCP renegotiation:

- [26-164] Ipv4-release-control
- [55] Event-Timestamp
- [25] Class

- [97] Framed-Ipv6-Prefix
- [123] Delegated-Ipv6-Prefix
- [44] Acct-Session-Id

The other attributes supported in Interim-Acct messages are included, based on the user configuration settings.

The following RADIUS attributes are never included in Interim-Acct messages when IPv4 address assignment fails after an IPCP renegotiation:

- [8] Framed-Ip-Address
- [9] Framed-Ip-Netmask

Interim-Acct Messages When an IPv4 Address is Released

The following RADIUS attributes are always included in Interim-Acct messages when an IPv4 address is released upon the termination of a subscriber session:

- [26-164] Ipv4-release-control
- [55] Event-Timestamp
- [25] Class
- [97] Framed-Ipv6-Prefix
- [123] Delegated-Ipv6-Prefix
- [44] Acct-Session-Id

The other attributes supported in Interim-Acct messages are included, based on the user configuration settings.

The following RADIUS attributes are never included in Interim-Acct messages when an IPv4 address is released to the AAA server:

- [8] Framed-Ip-Address
- [9] Framed-Ip-Netmask

For both these attributes, even if the IP address assigned to the user is 0.0.0.0, it is discarded and not included in the accounting messages.

Related Documentation

- [AAA Access Messages During IPCP Negotiations for Dual-Stack Subscribers on page 201](#)
- [Juniper Networks VSAs on page 248](#)
- [Juniper Networks VSAs Supported for Subscriber AAA Accounting Messages on page 196](#)
- *IPCP Renegotiation of IPv4 Addresses for Dual-Stack Subscribers*

DSL Forum VSAs in AAA Access and Accounting Messages Overview

JunosE Software supports the inclusion of a set of DSL Forum vendor-specific attributes (VSAs) in the following AAA access and accounting messages:

- Access-Request
- Acct-Start
- Acct-Stop
- Interim-Acct (if Acct-Stop messages are specified)
- COA-Request

The DSL Forum VSAs convey information about the subscriber associated with the digital subscriber line (DSL) and the data rate of the DSL. When you use **radius include dsl-forum-attributes** command to enable inclusion of the DSL Forum VSAs in these AAA messages, the router includes all of the attributes listed in [Table 48 on page 206](#) in the specified message, provided that the VSA is available in the information that the router receives from the digital subscriber line access multiplexer (DSLAM).



NOTE: JunosE Software also supports several Juniper Networks VSAs that you can use to include DSL-related information. See [“Juniper Networks VSAs” on page 248](#).

Related Documentation

- [Subscriber AAA Access Messages Overview on page 182](#)
- [Subscriber AAA Accounting Messages Overview on page 192](#)
- [DSL Forum VSAs Supported for AAA Access and Accounting Messages on page 206](#)
- [CLI Commands Used to Include DSL Forum VSAs in Access and Accounting Messages on page 217](#)
- [DSL Forum VSAs on page 260](#)

DSL Forum VSAs Supported for AAA Access and Accounting Messages

[Table 48 on page 206](#) lists the DSL Forum VSAs supported by JunosE Software in Access-Request, Acct-Start, Acct-Stop, (if Acct-Stop is specified) Interim-Acct, and COA-Request messages. JunosE Software uses the vendor ID assigned to the DSL Forum (3561, or DE9 in hexadecimal format) by the IANA.

Table 48: DSL Forum (Vendor ID 3561) VSAs Supported in AAA Access and Accounting Messages

Attribute Number	Attribute Name	Access-Request	Acct-Start	Acct-Stop	Interim-Acct	COA-Request
[26-1]	Agent-Circuit-Id	✓	✓	✓	✓	✓

Table 48: DSL Forum (Vendor ID 3561) VSAs Supported in AAA Access and Accounting Messages *(continued)*

Attribute Number	Attribute Name	Access-Request	Acct-Start	Acct-Stop	Interim-Acct	COA-Request
[26-2]	Agent-Remote-Id	✓	✓	✓	✓	✓
[26-129]	Actual-Data-Rate-Upstream	✓	✓	✓	✓	–
[26-130]	Actual-Data-Rate-Downstream	✓	✓	✓	✓	–
[26-131]	Minimum-Data-Rate-Upstream	✓	✓	✓	✓	–
[26-132]	Minimum-Data-Rate-Downstream	✓	✓	✓	✓	–
[26-133]	Attainable-Data-Rate-Upstream	✓	✓	✓	✓	–
[26-134]	Attainable-Data-Rate-Downstream	✓	✓	✓	✓	–
[26-135]	Maximum-Data-Rate-Upstream	✓	✓	✓	✓	–
[26-136]	Maximum-Data-Rate-Downstream	✓	✓	✓	✓	–
[26-137]	Minimum-Data-Rate-Upstream-Low-Power	✓	✓	✓	✓	–
[26-138]	Minimum-Data-Rate-Downstream-Low-Power	✓	✓	✓	✓	–
[26-139]	Maximum-Interleaving-Delay-Upstream	✓	✓	✓	✓	–
[26-140]	Actual-Interleaving-Delay-Upstream	✓	✓	✓	✓	–
[26-141]	Maximum-Interleaving-Delay-Downstream	✓	✓	✓	✓	–
[26-142]	Actual-Interleaving-Delay-Downstream	✓	✓	✓	✓	–
[26-144]	Access-Loop-Encapsulation	✓	✓	✓	✓	–
[26-254]	IWF-Session	✓	✓	✓	✓	–

Related Documentation

- [Subscriber AAA Access Messages Overview on page 182](#)
- [Subscriber AAA Accounting Messages Overview on page 192](#)
- [DSL Forum VSAs in AAA Access and Accounting Messages Overview on page 206](#)
- [CLI Commands Used to Include DSL Forum VSAs in Access and Accounting Messages on page 217](#)
- [DSL Forum VSAs on page 260](#)

RADIUS Attributes Supported for CLI AAA Messages

There are four types of AAA messages used by CLI users to gain administrative access to the router. Access-Challenge attributes pertain only to CLI/telnet users.

- Access-Request
- Access-Accept
- Access-Challenge
- Access-Reject

Table 49 on page 208 lists the RADIUS attributes supported for CLI AAA messages.

Table 49: CLI AAA Access Message RADIUS Attributes Supported

Attribute Number	Attribute Name	Access-Request	Access-Accept	Access-Challenge	Access-Reject
[1]	User-Name	✓	–	–	–
[2]	User Password	✓	–	–	–
[4]	NAS-IP-Address	✓	–	–	–
[6]	Service-Type	✓	✓	–	–
[18]	Reply-Message	–	–	✓	✓
[24]	State (Access-Request is only in response to an Access-Challenge)	✓	–	✓	–
[25]	Class	–	✓	–	–
[26-1]	Virtual-Router	–	✓	–	–
[26-18]	Init-CLI-Access-Level	–	✓	–	–
[26-19]	Allow-All-VR-Access	–	✓	–	–
[26-20]	Alt-CLI-Access-Level	–	✓	–	–
[26-21]	Alt-CLI-Virtual-Router-Name	–	✓	–	–
[26-25]	Redirect-Vrouter-Name	–	✓	–	–

- Related Documentation**
- [Subscriber AAA Access Messages Overview on page 182](#)
 - [Subscriber AAA Accounting Messages Overview on page 192](#)

- [CLI Commands Used to Configure RADIUS IETF Attributes on page 209](#)
- [CLI Commands Used to Configure Juniper Networks VSAs on page 213](#)
- [CLI Commands Used to Include or Exclude Attributes in RADIUS Messages on page 218](#)
- [CLI Commands Used to Ignore Attributes when Receiving Access-Accept Messages on page 223](#)
- [DSL Forum VSAs on page 260](#)
- [Juniper Networks VSAs on page 248](#)

CLI Commands Used to Modify RADIUS Attributes

You can configure the RADIUS Internet Engineering Task Force (IETF) attributes and the Juniper Networks vendor-specific attributes using CLI commands.

For many attributes, you can configure the router to include the attribute in RADIUS messages.

You can also configure the router to ignore many attributes that it receives in Access-Accept messages.

For a complete list of RADIUS attributes supported by JunosE Software, see [“RADIUS IETF Attributes” on page 241](#).

Related Documentation

- [CLI Commands Used to Configure RADIUS IETF Attributes on page 209](#)
- [CLI Commands Used to Configure Juniper Networks VSAs on page 213](#)
- [CLI Commands Used to Include ANCP-Related Juniper Networks VSAs in Access and Accounting Messages on page 215](#)
- [CLI Commands Used to Include DSL Forum VSAs in Access and Accounting Messages on page 217](#)
- [CLI Commands Used to Include or Exclude Attributes in RADIUS Messages on page 218](#)
- [CLI Commands Used to Ignore Attributes when Receiving Access-Accept Messages on page 223](#)

CLI Commands Used to Configure RADIUS IETF Attributes

[Table 50 on page 210](#) lists the RADIUS IETF attributes and the corresponding CLI commands used to configure them. The attributes are listed numerically—each attribute is followed by a list of the commands that you can use to manage the attribute.

Table 50: CLI Commands Used to Configure RADIUS IETF Attributes

Attribute Number	Attribute Name	CLI Command
[4]	NAS-IP-Address	<ul style="list-style-type: none"> <i>radius override nas-ip-addr tunnel-client-endpoint</i> <i>radius override nas-info</i>
[5]	NAS-Port	<ul style="list-style-type: none"> <i>radius include nas-port</i> <i>radius nas-port-format</i> <i>radius nas-port-format extended atm</i> <i>radius nas-port-format extended ethernet</i> <i>radius pppoe nas-port-format unique</i> <i>radius vlan nas-port-format stacked</i>
[8]	Framed-IP-Address	<ul style="list-style-type: none"> <i>radius include framed-ip-addr</i>
[9]	Framed-Ip-Netmask	<ul style="list-style-type: none"> <i>radius include framed-ip-netmask</i> <i>radius ignore framed-ip-netmask</i>
[13]	Framed-Compression	<ul style="list-style-type: none"> <i>radius include framed-compression</i>
[22]	Framed-Route	<ul style="list-style-type: none"> <i>radius include framed-route</i>
[25]	Class	<ul style="list-style-type: none"> <i>radius include class</i>
[30]	Called-Station-Id	<ul style="list-style-type: none"> <i>radius include called-station-id</i>
[31]	Calling-Station-Id	<ul style="list-style-type: none"> <i>radius calling-station-format</i> <i>radius calling-station-delimiter</i> <i>radius include calling-station-id</i> <i>radius override calling-station-id remote-circuit-id</i>
[32]	NAS-Identifier	<ul style="list-style-type: none"> <i>radius nas-identifier</i> <i>radius include nas-identifier</i> <i>radius override nas-info</i> <i>radius remote-circuit-id-format</i> <i>radius remote-circuit-id-delimiter</i>
[41]	Acct-Delay-Time	<ul style="list-style-type: none"> <i>radius include acct-delay-time</i>
[44]	Acct-Session-Id	<ul style="list-style-type: none"> <i>radius include acct-session-id</i> <i>radius acct-session-id-format</i>
[45]	Acct-Authentic	<ul style="list-style-type: none"> <i>radius include acct-authentic</i>
[49]	Acct-Terminate-Cause	<ul style="list-style-type: none"> <i>radius include acct-terminate-cause</i>
[50]	Acct-Multi-Session-Id	<ul style="list-style-type: none"> <i>radius include acct-multi-session-id</i>

Table 50: CLI Commands Used to Configure RADIUS IETF Attributes (*continued*)

Attribute Number	Attribute Name	CLI Command
[51]	Acct-Link-Count	• <i>radius include acct-link-count</i>
[52]	Acct-Input-Gigawords	• <i>radius include input-gigawords</i>
[53]	Output-Gigawords	• <i>radius include output-gigawords</i>
[55]	Event-Timestamp	• <i>radius include event-timestamp</i>
[61]	NAS-Port-Type	• <i>radius dsl-port-type</i> • <i>radius ethernet-port-type</i> • <i>radius include nas-port-type</i>
[64]	Tunnel-Type	• <i>radius include tunnel-type</i>
[65]	Tunnel-Medium-Type	• <i>radius include tunnel-medium-type</i>
[66]	Tunnel-Client-Endpoint	• <i>radius include tunnel-client-endpoint</i>
[67]	Tunnel-Server-Endpoint	• <i>radius include tunnel-server-endpoint</i>
[68]	Acct-Tunnel-Connection	• <i>radius include acct-tunnel-connection</i>
[77]	Connect-Info	• <i>radius connect-info-format l2tp-connect-speed</i> • <i>radius include connect-info</i>
[82]	Tunnel-Assignment-Id	• <i>radius include tunnel-assignment-id</i>
[83]	Tunnel-Preference	• <i>radius include tunnel-preference</i>
[87]	NAS-Port-Id	• <i>aaa intf-desc-format include</i> • <i>radius include nas-port-id</i> • <i>radius override nas-port-id remote-circuit-id</i>
[90]	Tunnel-Client-Auth-Id	• <i>radius include tunnel-client-auth-id</i>
[91]	Tunnel-Server-Auth-Id	• <i>radius include tunnel-server-auth-id</i>
[96]	Framed-Interface-Id	• <i>radius include framed-interface-id</i>
[97]	Framed-Ipv6-Prefix	• <i>radius include framed-ipv6-prefix</i>
[99]	Framed-Ipv6-Route	• <i>radius include framed-ipv6-route</i>
[100]	Framed-Ipv6-Pool	• <i>radius include framed-ipv6-pool</i>

Table 50: CLI Commands Used to Configure RADIUS IETF Attributes (*continued*)

Attribute Number	Attribute Name	CLI Command
[123]	Delegated-Ipv6-Prefix	• <i>radius include delegated-ipv6-prefix</i>
[144]	DS-Lite-Tunnel-Name	• <i>radius include ds-lite-tunnel-name</i>
[188]	Ascend-Num-In-Multilink	• <i>radius include ascend-num-in-multilink</i>
	All Tunnel Server Attributes	• <i>radius include tunnel-server-attributes</i>

Related Documentation

- [Propagation of LAG Subscriber Information to AAA and RADIUS on page 45](#)
- [RADIUS IETF Attributes Supported for Subscriber AAA Access Messages on page 183](#)
- [RADIUS IETF Attributes Supported for Subscriber AAA Accounting Messages on page 193](#)
- [RADIUS IETF Attributes Supported for AAA Tunnel Accounting Messages on page 200](#)
- [RADIUS IETF Attributes on page 241](#)
- *aaa intf-desc-format include*
- *radius acct-session-id-format*
- *radius calling-station-delimiter*
- *radius calling-station-format*
- *radius connect-info-format*
- *radius dsl-port-type*
- *radius ethernet-port-type*
- *radius ignore*
- *radius include*
- *radius nas-identifier*
- *radius nas-port-format*
- *radius nas-port-format extended*
- *radius override calling-station-id remote-circuit-id*
- *radius override nas-info*
- *radius override nas-ip-addr tunnel-client-endpoint*
- *radius override nas-port-id remote-circuit-id*
- *radius pppoe nas-port-format unique*
- *radius remote-circuit-id-delimiter*
- *radius remote-circuit-id-format*

- *radius vlan nas-port-format stacked*

CLI Commands Used to Configure Juniper Networks VSAs

Table 51 on page 213 lists the Juniper Networks VSAs and the corresponding CLI commands used to modify them. The attributes are listed numerically.

Table 51: CLI Commands Used to Configure Juniper Networks VSAs

Attribute Number	Attribute Name	CLI Command
[26-1]	Virtual-Router	• <i>radius ignore virtual-router</i>
[26-10]	Ingress-Policy-Name	• <i>radius include ingress-policy-name</i> • <i>radius ignore ingress-policy-name</i>
[26-11]	Egress-Policy-Name	• <i>radius include egress-policy-name</i> • <i>radius ignore egress-policy-name</i>
[26-14]	Service-Category	• <i>radius ignore atm-service-category</i>
[26-15]	PCR	• <i>radius ignore atm-pcr</i>
[26-16]	SCR	• <i>radius ignore atm-scr</i>
[26-17]	MBS	• <i>radius ignore atm-mbs</i>
[26-24]	Pppoe-Description	• <i>radius include pppoe-description</i>
[26-26]	QoS-Profile-Name	• <i>radius include qos-profile-name</i>
[26-35]	Acct-Input-Gigapackets	• <i>radius include input-gigapkts</i>
[26-36]	Acct-Output-Gigapackets	• <i>radius include output-gigapkts</i>
[26-44]	Tunnel-Interface-Id	• <i>radius include tunnel-interface-id</i>
[26-45]	Ipv6-Virtual-Router	• <i>radius include ipv6-virtual-router</i>
[26-46]	Ipv6-Local-Interface	• <i>radius include ipv6-local-interface</i>
[26-47]	Ipv6-Primary-DNS	• <i>radius include ipv6-primary-dns</i>
[26-48]	Ipv6-Secondary-DNS	• <i>radius include ipv6-secondary-dns</i>
[26-51]	Disconnect-Cause	• <i>radius include</i> <i>l2tp-ppp-disconnect-cause</i>
[26-53]	Service-Description	• <i>radius include</i> <i>profile-service-description</i>

Table 51: CLI Commands Used to Configure Juniper Networks VSAs (*continued*)

Attribute Number	Attribute Name	CLI Command
[26-55]	DHCP-Options	<ul style="list-style-type: none"> • <i>radius include</i> dhcp-options
[26-56]	DHCP-MAC-Address	<ul style="list-style-type: none"> • <i>radius include</i> dhcp-mac-address
[26-57]	DHCP-GI-Address	<ul style="list-style-type: none"> • <i>radius include</i> dhcp-gi-address
[26-62]	MLPPP-Bundle-Name	<ul style="list-style-type: none"> • <i>radius include</i> mlppp-bundle-name
[26-63]	Interface-Desc	<ul style="list-style-type: none"> • <i>radius include</i> interface-description
[26-81]	L2C-Information	<ul style="list-style-type: none"> • <i>radius include</i> access-loop-parameters
[26-92]	L2C-Up-Stream-Data	<ul style="list-style-type: none"> • <i>radius include</i> l2c-upstream-data
[26-93]	L2C-Down-Stream-Data	<ul style="list-style-type: none"> • <i>radius include</i> l2c-downstream-data
[26-106]	Ipv6-Ingress-Policy-Name	<ul style="list-style-type: none"> • <i>radius include</i> ipv6-ingress-policy-name • <i>radius ignore</i> ipv6-ingress-policy-name
[26-107]	Ipv6-Egress-Policy-Name	<ul style="list-style-type: none"> • <i>radius include</i> ipv6-egress-policy-name • <i>radius ignore</i> ipv6-egress-policy-name
[26-129]	Ipv6-NdRa-Prefix	<ul style="list-style-type: none"> • <i>radius include</i> ipv6-nd-ra-prefix
[26-141]	Downstream-Calculated-Qos-Rate	<ul style="list-style-type: none"> • <i>radius include</i> downstream-calculated-qos-rate access-request • <i>radius include</i> downstream-calculated-qos-rate acct-start • <i>radius include</i> downstream-calculated-qos-rate acct-stop
[26-142]	Upstream-Calculated-Qos-Rate	<ul style="list-style-type: none"> • <i>radius include</i> upstream-calculated-qos-rate access-request • <i>radius include</i> upstream-calculated-qos-rate acct-start • <i>radius include</i> upstream-calculated-qos-rate acct-stop
[26-143]	Max-Clients-Per-Interface	<ul style="list-style-type: none"> • <i>radius ignore</i> pppoe-max-session

Table 51: CLI Commands Used to Configure Juniper Networks VSAs (*continued*)

Attribute Number	Attribute Name	CLI Command
[26-150]	ICR-Partition-Id	<ul style="list-style-type: none"> • <i>radius include icr-partition-id</i> • <i>radius icr-partition-accounting</i>
[26-151]	IPv6-Acct-Input-Octets	<ul style="list-style-type: none"> • <i>radius include ipv6-accounting</i>
[26-152]	IPv6-Acct-Output-Octets	<ul style="list-style-type: none"> • <i>radius include ipv6-accounting</i>
[26-153]	IPv6-Acct-Input-Packets	<ul style="list-style-type: none"> • <i>radius include ipv6-accounting</i>
[26-154]	IPv6-Acct-Output-Packets	<ul style="list-style-type: none"> • <i>radius include ipv6-accounting</i>
[26-155]	IPv6-Acct-Input-Gigawords	<ul style="list-style-type: none"> • <i>radius include ipv6-accounting</i>
[26-156]	IPv6-Acct-Output-Gigawords	<ul style="list-style-type: none"> • <i>radius include ipv6-accounting</i>
[26-159]	DHCP-Option 82	<ul style="list-style-type: none"> • <i>radius include dhcp-option-82</i>
[26-165]	PCP-Server-Name	<ul style="list-style-type: none"> • <i>radius include pcp-server-name</i>

Related Documentation

- [Juniper Networks VSAs Supported for Subscriber AAA Access Messages on page 186](#)
- [Juniper Networks VSAs Supported for Subscriber AAA Accounting Messages on page 196](#)
- [RADIUS Attributes Supported for CLI AAA Messages on page 208](#)
- [Juniper Networks VSAs on page 248](#)
- [Monitoring Included RADIUS Attributes on page 295](#)
- [Monitoring Ignored RADIUS Attributes on page 298](#)
- *radius icr-partition-accounting*
- *radius ignore*
- *radius include*

CLI Commands Used to Include ANCP-Related Juniper Networks VSAs in Access and Accounting Messages

You use the **radius include** command to specify information about ANCP, also known as L2C, that you want to include in the RADIUS Access-Request, Acct-Start, and Acct-Stop messages. Also, if you specify Acct-Stop messages, the router includes ANCP information in Interim-Acct messages that the router sends to RADIUS. By default, the router does not include the ANCP-related information provided by the Juniper Networks VSAs in RADIUS messages.

These Juniper Networks ANCP-related VSAs are based on definitions in GSMP extensions for layer2 control (L2C) Topology Discovery and Line

Configuration—draft-wadhwa-gsmp-l2control-configuration-00.txt (July 2006 expiration).



NOTE:

- You must enable ANCP discovery with the **discovery-mode** command prior to configuring the **radius include** command with the ANCP-related VSAs. Configuring discovery mode enables the RADIUS authentication server to retrieve ANCP information.
- JunosE Software continues to support DSL Forum VSAs (vendor ID 3561) that you can use to include DSL-related information in RADIUS messages. See [“DSL Forum VSAs” on page 260](#).

Table 52 on page 216 lists the ANCP (L2C)-related keywords that you can use in the **radius include** command and the associated Juniper Networks VSAs. The table also indicates the mappings between ANCP parameters and the VSAs.

Table 52: ANCP (L2C)-Related Keywords for radius include Command

Command Keyword	Juniper Networks VSA Number	Juniper Networks VSA Name	ANCP Type	ANCP Subtype
l2cd-acc-loop-cir-id	[26-110]	Acc-Loop-Cir-Id	1	–
l2cd-acc-aggr-cir-id-bin	[26-111]	Acc-Aggr-Cir-Id-Bin	2	–
l2cd-acc-aggr-cir-id-asc	[26-112]	Acc-Aggr-Cir-Id-Asc	3	–
l2cd-act-data-rate-up	[26-113]	Act-Data-Rate-Up	4	129
l2cd-act-data-rate-dn	[26-114]	Act-Data-Rate-Dn	4	130
l2cd-min-data-rate-up	[26-115]	Min-Data-Rate-Up	4	131
l2cd-min-data-rate-dn	[26-116]	Min-Data-Rate-Dn	4	132
l2cd-att-data-rate-up	[26-117]	Att-Data-Rate-Up	4	133
l2cd-att-data-rate-dn	[26-118]	Att-Data-Rate-Dn	4	134
l2cd-max-data-rate-up	[26-119]	Max-Data-Rate-Up	4	135
l2cd-max-data-rate-dn	[26-120]	Max-Data-Rate-Dn	4	136
l2cd-min-lp-data-rate-up	[26-121]	Min-LP-Data-Rate-Up	4	137
l2cd-min-lp-data-rate-dn	[26-122]	Min-LP-Data-Rate-Dn	4	138
l2cd-max-interlv-delay-up	[26-123]	Max-Interlv-Delay-Up	4	139

Table 52: ANCP (L2C)-Related Keywords for radius include Command (continued)

Command Keyword	Juniper Networks VSA Number	Juniper Networks VSA Name	ANCP Type	ANCP Subtype
l2cd-act-interlv-delay-up	[26-124]	Act-Interlv-Delay-Up	4	140
l2cd-max-interlv-delay-dn	[26-125]	Max-Interlv-Delay-Dn	4	141
l2cd-act-interlv-delay-dn	[26-126]	Act-Interlv-Delay-Dn	4	142
l2cd-dsl-line-state	[26-127]	DSL-Line-State	4	143
l2cd-dsl-type	[26-128]	DSL-Type	4	144

Related Documentation

- [Subscriber AAA Access Messages Overview on page 182](#)
- [Subscriber AAA Accounting Messages Overview on page 192](#)
- [Juniper Networks VSAs Supported for Subscriber AAA Access Messages on page 186](#)
- [Juniper Networks VSAs Supported for Subscriber AAA Accounting Messages on page 196](#)
- [CLI Commands Used to Include or Exclude Attributes in RADIUS Messages on page 218](#)
- [Juniper Networks VSAs on page 248](#)
- [Monitoring Included RADIUS Attributes on page 295](#)
- [radius include](#)

CLI Commands Used to Include DSL Forum VSAs in Access and Accounting Messages

You can use the **radius include dsl-forum-attributes** command to control the inclusion of a set of DSL Forum VSAs in Access-Request, Acct-Start, Acct-Stop, and (if Acct-Stop messages are specified) Interim-Acct messages that the router sends to RADIUS.

The DSL Forum VSAs, as defined in RFC 4679—DSL Forum Vendor-Specific RADIUS Attributes (September 2006), convey information about the associated subscriber for and data rate of the DSL. A service provider might find it useful to enable inclusion of the DSL Forum VSAs in RADIUS messages in order to bill subscribers for different classes of service based on the data rate of their DSL connection.



NOTE: JunosE Software also supports several Juniper Networks VSAs that you can use to include DSL-related information. See [“Juniper Networks VSAs” on page 248](#).

The router receives data containing one or more of the DSL Forum VSAs from a DSLAM connected to the router via a PPPoE interface. When you enable the inclusion of the DSL Forum VSAs in these RADIUS messages, the router includes all of the following attributes

in the specified message type, provided that the VSA is available in the information that the router receives from the DSLAM.



NOTE: The router uses the vendor ID assigned to the DSL Forum (3561, or DE9 in hexadecimal format) by the IANA for the DSL Forum VSAs.

Agent-Circuit-Id [26-1]	Maximum-Data-Rate-Downstream [26-136]
Agent-Remote-Id [26-2]	Minimum-Data-Rate-Upstream-Low-Power [26-137]
Actual-Data-Rate-Upstream [26-129]	Minimum-Data-Rate-Downstream-Low-Power [26-138]
Actual-Data-Rate-Downstream [26-130]	Maximum-Interleaving-Delay-Upstream [26-139]
Minimum-Data-Rate-Upstream [26-131]	Actual-Interleaving-Delay-Upstream [26-140]
Minimum-Data-Rate-Downstream [26-132]	Maximum-Interleaving-Delay-Downstream [26-141]
Attainable-Data-Rate-Upstream [26-133]	Actual-Interleaving-Delay-Downstream [26-142]
Attainable-Data-Rate-Downstream [26-134]	Access-Loop-Encapsulation [26-144]
Maximum-Data-Rate-Upstream [26-135]	IWF-Session [26-254]

For information about enabling the QoS downstream rate application to obtain downstream rates from the Actual-Data-Rate-Downstream [26-130] DSL Forum VSA, see the *Configuring the Downstream Rate Using QoS Parameters* chapter in the *JunosE Quality of Service Configuration Guide*.

Related Documentation

- [Subscriber AAA Access Messages Overview on page 182](#)
- [Subscriber AAA Accounting Messages Overview on page 192](#)
- [DSL Forum VSAs in AAA Access and Accounting Messages Overview on page 206](#)
- [DSL Forum VSAs Supported for AAA Access and Accounting Messages on page 206](#)
- [DSL Forum VSAs on page 260](#)
- *radius include dsl-forum-attributes*

CLI Commands Used to Include or Exclude Attributes in RADIUS Messages

You can use the **radius include** command to enable or disable the inclusion of RADIUS attributes in Acct-on, Acct-off, Access-Request, Acct-Start, and Acct-Stop messages.

Table 53 on page 219 lists the RADIUS attributes that can be included or excluded in RADIUS messages using the **radius include** command and the RADIUS messages in which the attributes are supported.

Table 53: RADIUS Attributes Included in Corresponding RADIUS Messages

Attribute Number	Attribute Name	Access-Request	Acct-on	Acct-off	Acct-Start	Acct-Stop
[5]	NAS-Port	✓	–	–	✓	✓
[8]	Framed-IP-Address	✓	–	–	✓	✓
[9]	Framed-IP-Netmask	–	–	–	✓	✓
[13]	Framed-Compression	–	–	–	✓	✓
[22]	Framed-Route	–	–	–	✓	✓
[25]	Class	–	–	–	✓	✓
[26-10]	Ingress-Policy-Name	–	–	–	✓	✓
[26-11]	Egress-Policy-Name	–	–	–	✓	✓
[26-24]	Pppoe-Description	✓	–	–	✓	✓
[26-26]	QoS-Profile-Name	–	–	–	✓	✓
[26-35]	Acct-Input-Gigapackets	–	–	–	–	✓
[26-43]	Acct-Output-Gigapackets	–	–	–	–	✓
[26-44]	Tunnel-Interface-ID	✓	–	–	✓	✓
[26-45]	Ipv6-Virtual-Router	–	–	–	–	✓
[26-46]	Ipv6-Local-Interface	–	–	–	–	✓
[26-47]	Ipv6-Primary-DNS	–	–	–	–	✓
[26-48]	Ipv6-Secondary-DNS	–	–	–	–	✓
[26-51]	Disconnect-Cause	–	–	–	–	✓
[26-53]	Service-Description	✓	–	–	✓	✓
[26-55]	DHCP-Options	✓	–	–	✓	✓
[26-56]	DHCP-MAC-Address	✓	–	–	✓	✓

Table 53: RADIUS Attributes Included in Corresponding RADIUS Messages (*continued*)

Attribute Number	Attribute Name	Access-Request	Acct-on	Acct-off	Acct-Start	Acct-Stop
[26-57]	DHCP-GI-Address	✓	–	–	✓	✓
[26-62]	MLPPP-Bundle-Name	✓	–	–	✓	✓
[26-63]	Interface-Description	✓	–	–	✓	✓
[26-81]	L2c-Information	✓	–	–	–	–
[26-92]	L2C-Up-Stream-Data	✓	–	–	✓	✓
[26-93]	L2C-Down-Stream-Data	✓	–	–	✓	✓
[26-106]	Ipv6-Ingress-Policy-Name	–	–	–	✓	✓
[26-107]	Ipv6-Egress-Policy-Name	–	–	–	✓	✓
[26-110]	Acc-Loop-Cir-Id	✓	–	–	✓	✓
[26-111]	Acc-Aggr-Cir-Id-Bin	✓	–	–	✓	✓
[26-112]	Acc-Aggr-Cir-Id-Asc	✓	–	–	✓	✓
[26-113]	Act-Data-Rate-Up	✓	–	–	✓	✓
[26-114]	Act-Data-Rate-Dn	✓	–	–	✓	✓
[26-115]	Min-Data-Rate-Up	✓	–	–	✓	✓
[26-116]	Min-Data-Rate-Dn	✓	–	–	✓	✓
[26-117]	Att-Data-Rate-Up	✓	–	–	✓	✓
[26-118]	Att-Data-Rate-Dn	✓	–	–	✓	✓
[26-119]	Max-Data-Rate-Up	✓	–	–	✓	✓
[26-120]	Max-Data-Rate-Dn	✓	–	–	✓	✓
[26-121]	Min-LP-Data-Rate-Up	✓	–	–	✓	✓
[26-122]	Min-LP-Data-Rate-Dn	✓	–	–	✓	✓
[26-123]	Max-Interlv-Delay-Up	✓	–	–	✓	✓
[26-124]	Act-Interlv-Delay-Up	✓	–	–	✓	✓

Table 53: RADIUS Attributes Included in Corresponding RADIUS Messages (*continued*)

Attribute Number	Attribute Name	Access-Request	Acct-on	Acct-off	Acct-Start	Acct-Stop
[26-125]	Max-Interlv-Delay-Dn	✓	–	–	✓	✓
[26-126]	Act-Interlv-Delay-Dn	✓	–	–	✓	✓
[26-127]	DSL-Line-State	✓	–	–	✓	✓
[26-128]	DSL-Type	✓	–	–	✓	✓
[26-129]	Ipv6-NdRa-Prefix	–	–	–	–	✓
[26-141]	Downstream-Calculated-Qos	✓	–	–	✓	✓
[26-142]	Upstream-Calculated-Qos-Rate	✓	–	–	✓	✓
[26-150]	ICR-Partition-Id	✓	–	–	✓	✓
[26-159]	DHCP-Option 82	✓	–	–	✓	✓
[26-165]	PCP-Server-Name	–	–	–	✓	✓
[30]	Called-Station-Id	✓	–	–	✓	✓
[31]	Calling-Station-Id	✓	–	–	✓	✓
[32]	NAS-Identifier	✓	✓	✓	✓	✓
[41]	Acct-Delay-Time	–	✓	✓	–	–
[44]	Acct-Session-Id	✓	✓	✓	–	–
[45]	Acct-Authentic	–	✓	✓	–	–
[49]	Acct-Terminate-Cause	–	–	✓	–	–
[50]	Acct-Multi-Session-Id	✓	–	–	✓	✓
[51]	Acct-Link-Count	–	–	–	✓	✓
[52]	Acct-Input-Gigawords	–	–	–	–	✓
[53]	Acct-Output-Gigawords	–	–	–	–	✓
[55]	Event-Timestamp	–	✓	✓	✓	✓
[61]	NAS-Port-Type	✓	–	–	✓	✓

Table 53: RADIUS Attributes Included in Corresponding RADIUS Messages (*continued*)

Attribute Number	Attribute Name	Access-Request	Acct-on	Acct-off	Acct-Start	Acct-Stop
[64]	Tunnel-Type	✓	–	–	✓	✓
[65]	Tunnel-Medium-Type	✓	–	–	✓	✓
[66]	Tunnel-Client-Endpoint	✓	–	–	✓	✓
[67]	Tunnel-Server-Endpoint	✓	–	–	✓	✓
[68]	Acct-Tunnel-Connection	✓	–	–	✓	✓
[77]	Connect-Info	✓	–	–	✓	✓
[82]	Tunnel-Assignment-Id	–	–	–	✓	✓
[83]	Tunnel-Preference	–	–	–	✓	✓
[87]	NAS-Port-Id	✓	–	–	✓	✓
[90]	Tunnel-Client-Auth-Id	✓	–	–	✓	✓
[91]	Tunnel-Server-Auth-Id	✓	–	–	✓	✓
[96]	Framed-Interface-Id	✓	–	–	✓	✓
[97]	Framed-Ipv6-Prefix	✓	–	–	✓	✓
[99]	Framed-Ipv6-Route	–	–	–	–	✓
[100]	Framed-IPv6-Pool	–	–	–	–	✓
[123]	Delegated-IPv6-Prefix	–	–	–	–	✓
[144]	DS-Lite-Tunnel-Name	–	–	–	✓	✓
[188]	Ascend-Num-In-Multilink	✓	–	–	✓	✓
	All Tunnel-Server-Attributes	✓	–	–	✓	✓
	All Ipv6-Accounting Attributes	–	–	–	–	✓

Related Documentation

- [Subscriber AAA Access Messages Overview on page 182](#)
- [Subscriber AAA Accounting Messages Overview on page 192](#)
- [RADIUS IETF Attributes on page 241](#)
- [Juniper Networks VSAs on page 248](#)

- [Monitoring Included RADIUS Attributes on page 295](#)
- *radius include*

CLI Commands Used to Ignore Attributes when Receiving Access-Accept Messages

You can use the **radius ignore** command to configure the router to ignore or accept a RADIUS attribute from the received Access-Accept messages.

The following attributes can be ignored or accepted using the **radius ignore** command:

- atm-mbs
- atm-pcr
- atm-scr
- atm-service-category
- egress-policy-name
- framed-ip-netmask
- ingress-policy-name
- ipv6-egress-policy-name
- ipv6-ingress-policy-name
- pppoe-max-session
- virtual-router

Related Documentation

- [Subscriber AAA Access Messages Overview on page 182](#)
- [Monitoring Ignored RADIUS Attributes on page 298](#)
- *radius ignore*

RADIUS Per-Profile Attribute List Configuration Overview

JunosE Software enables you to configure RADIUS-specific attributes for subscribers attached to a specific PPP profile. If a per-profile list is configured, then only the attributes specified in the per-profile list are processed. If the per-profile list is not configured, then the existing standard attributes are configured.



NOTE: The attributes supported by the per-profile list take precedence over the standard RADIUS attribute configuration. By default, the inclusion of all attributes is disabled in the per-profile list.

This feature enables you to configure the following RADIUS attributes:

- **override nas-ip-addr**

- **calling-station-id**

- Related Documentation**
- [RADIUS Overview on page 180](#)
 - *attributes (RADIUS)*

Example: Configuring RADIUS-Specific Attributes

In this example, RADIUS-specific attributes are configured for subscribers attached to a specific PPP profile. You can configure this as follows:

1. Create a RADIUS per-profile attribute list, and configure the required RADIUS attributes in the list.

```
host1(config)#radius per-profile-attr-list abc
host1 (config-perprofile-list)#request-type acct-start
host1 (config-perprofile-list)#action-type enable
host1 (config-perprofile-list)#attributes calling-station-id override-nas-ip-addr
```

2. Create an AAA profile.

```
host1(config)#aaa profile aaaprofile1
```

3. Specify the RADIUS attribute list in the AAA profile.

```
host1(config-aaa-profile)#radius-perprofilelist-name abc
```

4. Create a PPP profile.

```
host1(config)#profile pppprofile1
```

5. Attach the AAA profile name to the PPP profile.

```
host1(config-profile)#ppp aaa-profile aaaprofile1
```

6. To view the attributes configured in the RADIUS per-profile attribute list, issue the **show radius per-profile-attr-list** command.

```
host1#show radius per-profile-attr-list abc
```

Attribute Name	AccessRequest	AccountStart	AccountStop
calling-station-id	enabled	disabled	enabled
override-nas-ip-addr	enabled	enabled	enabled

- Related Documentation**
- [RADIUS Per-Profile Attribute List Configuration Overview on page 223](#)
 - *aaa profile*
 - *action-type*
 - *attributes (RADIUS)*
 - *ppp aaa-profile*
 - *profile*
 - *radius per-profile-attr-list*
 - *radius-perprofilelist-name*

- *request-type*
- *show radius per-profile-attr-list*

CHAPTER 5

Configuring RADIUS Dynamic-Request Server

This chapter describes the RADIUS dynamic-request server feature on E Series routers. The following topics describe this feature:

- [RADIUS Dynamic-Request Server Overview on page 227](#)
- [RADIUS Dynamic-Request Server Platform Considerations on page 228](#)
- [RADIUS Dynamic-Request Server References on page 229](#)
- [Understanding RADIUS-Initiated Disconnect on page 229](#)
- [Configuring RADIUS-Initiated Disconnect on page 231](#)
- [Understanding RADIUS-Initiated Change of Authorization on page 232](#)
- [Configuring RADIUS-Initiated Change of Authorization on page 234](#)

RADIUS Dynamic-Request Server Overview

The E Series router's RADIUS dynamic-request server feature provides an efficient way for you to use RADIUS servers to centrally manage user sessions. The RADIUS dynamic-request server enables the router to receive the following types of messages from RADIUS servers:

- Disconnect messages—Immediately terminate specific user sessions.
- Change-of-Authorization (COA) messages—Dynamically modify session authorization attributes, such as data filters.




NOTE: The RADIUS dynamic-request server's support for COA messages is used by the Service Manager and by the E Series router's packet mirroring feature. For information about using the Service Manager, see the *Configuring Service Manager* chapter in this guide. For specific information about using the dynamic-request server with packet mirroring, see the *Configuring RADIUS-Based Packet Mirroring* chapter in the *JunosE Policy Management Configuration Guide*.

For example, you might use the RADIUS dynamic-request server to terminate specific user sessions. Without the RADIUS dynamic-request server, the only way to disconnect a RADIUS user is from the E Series router. This disconnect method is cumbersome when a network has many systems. The RADIUS dynamic-request server allows RADIUS servers to initiate user-related operations, such as a termination operation, by sending unsolicited request messages to an E Series router.

Figure 5 on page 228 shows a network that would benefit from the RADIUS dynamic-request server functionality. In Figure 5 on page 228, instead of disconnecting users on each E Series router, the RADIUS servers can initiate the disconnection. Although the network has multiple RADIUS servers, the servers share a common database that contains authorization and accounting information. Having a common database allows any server to view who is currently valid and connected, and allows service providers to manage the disconnection of users.

Figure 5: Sample Remote Access Network Using RADIUS



- Related Documentation**
- [Monitoring the Configuration of the RADIUS Dynamic-Request Server on page 300](#)
 - [Setting the Baseline for RADIUS Dynamic-Request Server Statistics on page 298](#)

RADIUS Dynamic-Request Server Platform Considerations

RADIUS dynamic-request server is supported on all E Series routers. For information about the modules supported on E Series routers:

- See the *ERX Module Guide* for modules supported on ERX7xx models, ERX14xx models, and the ERX310 Broadband Services Router.
- See the *E120 and E320 Module Guide* for modules supported on the E120 and E320 Broadband Services Routers.

- Related Documentation**
- [RADIUS Dynamic-Request Server Overview on page 227](#)

RADIUS Dynamic-Request Server References

For more information about the RADIUS dynamic-request server feature, see the following references:

- RFC 2865—Remote Authentication Dial In User Service (RADIUS) (June 2000)
- RFC 2866—RADIUS Accounting (June 2000)
- RFC 5176—Dynamic Authorization Extensions to Remote Authentication Dial In User Service (RADIUS) (January 2008)

**Related
Documentation**

- [RADIUS Dynamic-Request Server Overview on page 227](#)
- [RADIUS Dynamic-Request Server Platform Considerations on page 228](#)

Understanding RADIUS-Initiated Disconnect

In a typical client-server RADIUS environment, the E Series router functions as the client and the RADIUS server functions as the server. However, when using the RADIUS dynamic-request server feature, the roles are reversed. For example, during a RADIUS-initiated disconnect operation, the E Series router's RADIUS dynamic-request server functions as the server, and the RADIUS server functions as the disconnect client.

This section describes the RADIUS dynamic-request server's RADIUS-initiated disconnect feature.

Disconnect Messages

To centrally control the disconnection of remote access users, the RADIUS dynamic-request server on the router must receive and process unsolicited messages from RADIUS servers.

The RADIUS-initiated disconnect feature uses the existing format of RADIUS disconnect request and response messages. The RADIUS-initiated disconnect feature uses the following codes in its RADIUS request and response messages:

- Disconnect-Request (40)
- Disconnect-ACK (41)
- Disconnect-NAK (42)

Message Exchange

The RADIUS server and the router's RADIUS dynamic-request server exchange messages using User Datagram Protocol (UDP). The Disconnect-Request message sent by the RADIUS server has the same format as the COA-Request packet that is sent for a change of authorization operation.

The disconnect response is either a Disconnect-ACK or a Disconnect-NAK message:

- If AAA successfully disconnects the user, the response is a RADIUS-formatted packet with a Disconnect-ACK message.
- If AAA cannot disconnect the user, the request is malformed, or attributes are missing from the request, the response is a RADIUS-formatted packet with a Disconnect-NAK message.

Supported Error-Cause Codes (RADIUS Attribute 101)

When a disconnect request fails, the RADIUS dynamic-request server includes an error-cause attribute (RADIUS attribute 101) in the Disconnect-NAK message that it sends back to the RADIUS server. If the detected error does not map to one of the supported error-cause attributes, the router sends the Disconnect-NAK without an error-cause attribute. [Table 54 on page 230](#) lists the supported error-cause codes.

Table 54: Error-Cause Codes (RADIUS Attribute 101)

Code	Value	Description
401	Unsupported attribute	The request contains an attribute that is not supported (for example, a third-party attribute).
402	Missing attribute	A critical attribute (for example, the session identification attribute) is missing from a request.
404	Invalid request	Some other aspect of the request is invalid, such as if one or more attributes (for example, the packet mirroring Mirror Identifier value) are not formatted properly.
503	Session context not found	The session context identified in the request does not exist on the NAS.
504	Session context not removable	The subscriber identified by attributes in the disconnect request is owned by a component that does not support RADIUS-initiated disconnect (for example, IP LAC subscribers cannot be disconnected).
506	Resources unavailable	A request could not be honored due to lack of available NAS resources (such as memory).

Qualifications for Disconnect

For the server to disconnect a user, the Disconnect-Request message must contain an attribute with a session ID. The Disconnect-Request message can contain an Acct-Session-Id (44) attribute or an Acct-Multi-Session-Id (50) attribute for the session ID or both. If both the Acct-Session-Id and Acct-Multi-Session-Id attributes are present in the request, the router uses both attributes. If the User-Name (1) attribute is also present in the request, the username and session ID are used to perform the disconnection. Authentication, authorization, and accounting (AAA) services handle the actual request.



NOTE: The inclusion of the Acct-Multi-Session-Id (50) attribute in RADIUS Disconnect-Request messages for LAC L2TP sessions causes the disconnection of L2TP LAC user sessions to occur properly. The value of this attribute is constructed from the Acct-Session-ID (44) attribute of the first PPP link established for MLPPP bundles. If the Acct-Multi-Session-Id (50) attribute is contained in the Disconnect-Request message for MLPPP links, which are on the LAC side of an L2TP tunnel, the subscriber session is disconnected.

Security/Authentication

The RADIUS server (the disconnect client) must calculate the authenticator as specified for an Accounting-Request message in RFC 2866. The router's RADIUS dynamic-request server verifies the request using authenticator calculation as specified for an Accounting-Request message in RFC 2866. A key (secret), as specified in RFC 2865, must be configured and used in the calculation of the authenticator. The response authenticator is calculated as specified for an Accounting-Response message in RFC 2866.

Related Documentation

- [Configuring RADIUS-Initiated Disconnect on page 231](#)
- [Understanding RADIUS-Initiated Change of Authorization on page 232](#)
- [Configuring RADIUS-Initiated Change of Authorization on page 234](#)

Configuring RADIUS-Initiated Disconnect

To configure RADIUS-initiated disconnect feature, perform the following steps to set up the RADIUS dynamic-request server that will perform the disconnect operation:

1. Configure the RADIUS dynamic-request server, and enter RADIUS Configuration mode.


```
host1(config)#radius dynamic-request server 10.10.5.10
host1(config-radius)#
```
2. Enable the RADIUS-initiated disconnect capability on the RADIUS dynamic-request server.


```
host1(config-radius)#subscriber disconnect
```
3. Define the secret used in the RADIUS Authenticator field during exchanges between the RADIUS dynamic-request server and the RADIUS server.


```
host1(config-radius)#key Secret3Clientkey
```
4. (Optional) Specify the UDP port on which the RADIUS dynamic-request server listens for messages from the RADIUS server. The default is 1700.


```
host1(config-radius)#udp-port 1770
```

Related Documentation

- [Setting the Baseline for RADIUS Dynamic-Request Server Statistics on page 298](#)
- [Monitoring RADIUS Dynamic-Request Server Statistics on page 299](#)

- [Monitoring the Configuration of the RADIUS Dynamic-Request Server on page 300](#)
- *key*
- *radius disconnect client*
- *subscriber disconnect*
- *udp-port*

Understanding RADIUS-Initiated Change of Authorization

This section describes the RADIUS dynamic-request server's support for COA messages. COA messages are used by the E Series router's RADIUS-initiated packet mirroring feature, which is described in the *Configuring RADIUS-Based Packet Mirroring* chapter in the *JunosE Policy Management Configuration Guide*, and by Service Manager, which is described in the *Configuring Service Manager* chapter of this guide.

Change-of-Authorization Messages

The RADIUS dynamic-request server receives and processes the unsolicited COA messages from RADIUS servers. The RADIUS-initiated COA feature uses the following codes in its RADIUS request and response messages:

- COA-Request (43)
- COA-ACK (44)
- COA-NAK (45)

Message Exchange

The RADIUS server and the router's RADIUS dynamic-request server exchange messages using UDP. The COA-Request message sent by the RADIUS server has the same format as the Disconnect-Request packet that is sent for a disconnect operation.

The response is either a COA-ACK or a COA-NAK message:

- If AAA successfully changes the authorization, the response is a RADIUS-formatted packet with a COA-ACK message, and the data filter is applied to the session.
- If AAA is unsuccessful, the request is malformed, or attributes are missing, the response is a RADIUS-formatted packet with a COA-NAK message.

Supported Error-Cause Codes (RADIUS Attribute 101)

When AAA is unsuccessful, the RADIUS dynamic-request server includes an error-cause attribute (RADIUS attribute 101) in the COA-NAK message that it sends back to the RADIUS server. If the detected error does not map to one of the supported error-cause attributes, the router sends the COA-NAK without an error-cause attribute.

[Table 55 on page 233](#) lists the supported error-cause codes.

Table 55: Error-Cause Codes (RADIUS Attribute 101)

Code	Value	Description
401	Unsupported attribute	The request contains an attribute that is not supported (for example, a third-party attribute).
402	Missing attribute	A critical attribute (for example, the session identification attribute) is missing from a request.
404	Invalid request	Some other aspect of the request is invalid, such as if one or more attributes (for example, the packet mirroring Mirror Identifier value) are not formatted properly.
503	Session context not found	The session context identified in the request does not exist on the NAS.
504	Session context not removable	The subscriber identified by attributes in the disconnect request is owned by a component that does not support RADIUS-initiated disconnect (for example, IP LAC subscribers cannot be disconnected).
506	Resources unavailable	A request could not be honored due to lack of available NAS resources (such as memory).

Qualifications for Change of Authorization

To complete the change of authorization for a user, the COA-Request must contain one of the following RADIUS attributes or pairs of attributes. AAA services handle the actual request.

- User-Name [attribute 1] with Virtual-Router [attribute 26–1] to identify the user per virtual router context
- Framed-IP-Address [attribute 8] with Virtual-Router [attribute 26–1] to identify the address per virtual router context
- Calling-Station-ID [attribute 31]
- Acct-Session-ID [attribute 44] (mandatory for all COA requests, except when the request is for packet mirroring)
- Nas-Port-ID [attribute 5]
- DHCP-Option-82 [attribute 26–159], Vendor ID 4874
- Agent-Circuit-ID [attribute 26–1], Vendor ID 3561
- Agent-Remote-ID [attribute 26–2], Vendor ID 3561



NOTE: The Calling-Station-ID attribute is valid only for the tunneled subscribers and on the LNS. Additionally, the Calling-Station-ID and Nas-Port-ID attributes are valid only if there is no RADIUS override setting.

Security/Authentication

For change-of-authorization operations, the RADIUS server calculates the authenticator as specified for an Accounting-Request message in RFC 2866. The RADIUS dynamic-request server verifies the request using authenticator calculation as specified for an Accounting-Request in RFC 2866. A key (secret), as specified in RFC 2865, must be configured and used in the calculation of the authenticator. The response authenticator is calculated as specified for an Accounting-Response message in RFC 2866.

Related Documentation

- [Configuring RADIUS-Initiated Change of Authorization on page 234](#)
- [Understanding RADIUS-Initiated Disconnect on page 229](#)
- [Configuring RADIUS-Initiated Disconnect on page 231](#)

Configuring RADIUS-Initiated Change of Authorization

To configure the RADIUS dynamic-request change of authorization (COA) feature, perform the following steps to set up the RADIUS dynamic-request server that will perform the COA operation:

1. Configure the RADIUS dynamic-request server, and enter RADIUS Configuration mode.

```
host1(config)#radius dynamic-request server 10.10.5.10
```

2. Enable the COA capability on the RADIUS dynamic-request server.

```
host1(config-radius)#authorization change
```

3. Define the key (secret) used in the RADIUS Authenticator field during exchanges between the RADIUS dynamic-request server and the RADIUS server.

```
host1(config-radius)#key Secret21Clientkey
```

4. (Optional) Specify the UDP port on which the router listens for messages from the RADIUS server. The default is 1700.

```
host1(config-radius)#udp-port 1770
```

Related Documentation

- [Setting the Baseline for RADIUS Dynamic-Request Server Statistics on page 298](#)
- [Monitoring RADIUS Dynamic-Request Server Statistics on page 299](#)
- [Monitoring the Configuration of the RADIUS Dynamic-Request Server on page 300](#)
- *authorization change*
- *key*
- *udp-port*

CHAPTER 6

Configuring RADIUS Relay Server

This chapter describes the E Series router's RADIUS relay server feature. The RADIUS relay server provides authentication, authorization, accounting, and addressing services to wireless subscribers in public areas, such as airports and coffee shops. This chapter has the following sections:

- [Understanding the RADIUS Relay Server on page 235](#)
- [RADIUS Relay Server Platform Considerations on page 238](#)
- [RADIUS Relay Server References on page 238](#)
- [RADIUS Relay Server and the SRC Software on page 238](#)
- [Configuring RADIUS Relay Server Support on page 239](#)


Understanding the RADIUS Relay Server

The JunosE RADIUS relay server provides authentication, authorization, accounting, and addressing services in an 802.1x-based wireless environment.

The IEEE 802.1x standard is an authentication standard for wireless LANs; it enables a wireless subscriber to be authenticated by a central authority. The standard uses the Extensible Authentication Protocol (EAP) for message exchange during the authentication process. The E Series router's RADIUS relay server enhances the 802.1x environment by including authorization, accounting, and addressing support for wireless subscribers.

[Figure 6 on page 236](#) illustrates a typical 802.1x-based wireless environment. In the figure, wireless subscribers connect to wireless access points (WAPs) for authentication. The WAPs in turn connect to the E Series router's RADIUS relay server. The RADIUS relay server passes the request on to the authentication server, which might be a RADIUS or TACACS+ server. The RADIUS server authenticates the subscriber, who is then granted access. After authentication, the RADIUS relay server obtains an IP address for the subscriber from the Dynamic Host Configuration Protocol (DHCP) local or external server. The RADIUS relay server can also use the RADIUS server or the optional Session and Resource Control (SRC) software (formerly the SDX software), to provide the accounting support.

Figure 6: RADIUS Relay Server



How RADIUS Relay Server Works

When a wireless subscriber starts a session, the WAP encapsulates EAP attributes into a RADIUS Access-Request message and sends the request to the E Series router, which the WAP views as the RADIUS server. The encapsulated message uses the RADIUS EAP-Message (79) attribute. The RADIUS relay server does not process any of the EAP attributes in the RADIUS Access-Request message; the encrypted message is simply passed through the router to the actual RADIUS server. The RADIUS server must be EAP aware.

You can also use an optional RADIUS proxy server to provide additional enhancements to the 802.1x-based environment. For example, the RADIUS proxy server enables subscribers to be multiplexed to multiple Internet service providers (ISPs) that are customers of the same carrier. The server performs one of the following actions:

- If the ISP's RADIUS server supports EAP, the RADIUS proxy server extends the EAP session to the RADIUS server.
- If the ISP's RADIUS server does not support EAP, the RADIUS proxy server translates the EAP session into a legacy RADIUS session for the RADIUS server.

Authentication and Addressing

The WAP initiates the authentication and authorization request by sending a standard RADIUS Access-Request to the RADIUS relay server. The Access-Request must include the attributes listed in [Table 56 on page 236](#). The attributes uniquely identify the wireless subscriber.

Table 56: Required RADIUS Access-Request Attributes

Attribute Name	Description
Called-Station-id [30]	Subscriber's WAP
Calling-Station-id [31]	Subscriber's media access control (MAC) address

When the RADIUS server authenticates the subscriber, the router's RADIUS relay server creates a RADIUS Access-Accept message and sends the message back to the subscriber. The router's DHCP server (either the router's DHCP local server or an external DHCP server) assigns an IP address to the subscriber and creates the subscriber interface.

For information about using the optional SRC software with the RADIUS relay server to assign IP addresses, see the *Using the SRC Software for Addressing* section in “[RADIUS Relay Server and the SRC Software](#)” on page 238.

The WAP might periodically reauthenticate a subscriber. For example, reauthentication is necessary to renegotiate a new Wired Equivalent Privacy (WEP) key. The RADIUS relay server ignores any new RADIUS attributes that are sent during a renegotiation operation.

Accounting

The RADIUS relay server's clients (the WAPs) send standard accounting request messages to the RADIUS relay server. The accounting server processes the request and sends the results back to the RADIUS relay server, which then creates a RADIUS accounting response message and forwards the information to the client WAP.

For tracking purposes, the forwarding RADIUS relay server adds the Radius-Client-Address vendor-specific attribute (VSA 26-52) to the forwarded accounting request messages. The VSA indicates the RADIUS relay server's IP address.

For information about using the SRC software with the RADIUS relay server to provide accounting, see the *Using the SRC Software for Addressing* section in “[RADIUS Relay Server and the SRC Software](#)” on page 238.

[Table 57 on page 237](#) shows the RADIUS attributes that must be included in accounting requests. The attributes uniquely identify subscribers.

Table 57: Required RADIUS Accounting Attributes

For RADIUS Acct-Start and Acct-Stop Messages	Description
Called-Station-id [30]	Subscriber's WAP
Calling-Station-id [31]	Subscriber's MAC address
For RADIUS Acct-On and Acct-Off Messages	
Called-Station-id [30]	Subscriber's WAP

Terminating the Wireless Subscriber's Connection

The RADIUS relay server terminates the wireless subscriber's session when one of the following events occurs. When a subscriber session is terminated, the subscriber's IP address is released back into the available address pool.

- The RADIUS relay server receives a RADIUS accounting stop request.
- No RADIUS accounting messages are received for this subscriber for more than 24 hours.

- Related Documentation**
- [RADIUS Relay Server and the SRC Software on page 238](#)
 - [Configuring RADIUS Relay Server Support on page 239](#)

RADIUS Relay Server Platform Considerations

RADIUS relay is supported on all E Series routers.

For information about the modules supported on E Series routers:

- See the *ERX Module Guide* for modules supported on ERX7xx models, ERX14xx models, and the ERX310 Broadband Services Router.
- See the *E120 and E320 Module Guide* for modules supported on the E120 and E320 Broadband Services Routers.

- Related Documentation**
- [Understanding the RADIUS Relay Server on page 235](#)
 - [RADIUS Relay Server References on page 238](#)

RADIUS Relay Server References

For more information about RADIUS relay server, see the following resources:

- IEEE 802.1x-2001—Port-Based Network Access Control
- RFC 2869—RADIUS Extensions (June 2000)
- RFC 2284—PPP Extensible Authentication Protocol (EAP) (March 1998)
- RFC 3539—Authentication, Authorization and Accounting (AAA) Transport Profile (June 2003)

- Related Documentation**
- [Understanding the RADIUS Relay Server on page 235](#)
 - [RADIUS Relay Server Platform Considerations on page 238](#)

RADIUS Relay Server and the SRC Software

The SRC software is an advanced subscriber configuration and management service. The RADIUS relay server can optionally use the SRC software to perform addressing and accounting services for the subscriber and WAP.

The RADIUS relay server uses the E Series router's DHCP local server or DHCP external server and SRC client process to communicate with the SRC software.

Using the SRC Software for Addressing

If you integrate the SAE software into the RADIUS relay server configuration, the application can contribute to the address pool selection used to lease an address to the

subscriber. The SRC software only contributes to address pool selection when the DHCP local server is used; it is not supported when a DHCP external server is used.

Using the SRC Software for Accounting

If you use the SRC software with the RADIUS relay server feature, two accounting domains might actually be created. The first domain is established by the WAP, when the subscriber is authenticated. The second domain is created for the connection between the E Series router and the SRC software.

If you want to continue to use the SRC software's user session and problem-tracking features, you should *not* configure the SRC software to generate RADIUS accounting records. Also, the following attributes must be configured on the RADIUS server used by the WAP:

- Service-Bundle [26-31]
- Class [25]
- User-Name [1]

Related Documentation

- [Understanding the RADIUS Relay Server on page 235](#)

Configuring RADIUS Relay Server Support

To configure the RADIUS relay server feature, you enable support for the feature on the E Series router and identify the key (secret) used for the connection between the WAP and the RADIUS relay server. The following example configures a RADIUS relay authentication server. Use similar steps to configure a RADIUS relay accounting server.



NOTE: The E Series router supports one instance of the RADIUS relay server per virtual router. The instance can provide authentication, authorization, and accounting support.

1. Enable RADIUS relay server support on the E Series router, and enter RADIUS Relay Configuration mode.


```
host1(config)#radius relay authentication server
host1(config-radius-relay)#
```
2. Specify the IP address and mask of the network that will use the relay authentication server, and the secret used during exchanges between the relay authentication server and clients (the WAPs).


```
host1(config-radius-relay)#key 192.168.25.9 255.255.255.255 mysecret
```
3. Specify the router's User Datagram Protocol (UDP) port on which the RADIUS relay server listens.


```
host1(config-radius-relay)#udp-port 1812
```
4. (Optional) Verify the configuration.

```
host1(config-radius-relay)#exit
host1(config)#exit
host1#show radius relay servers
```

RADIUS Relay Authentication Server Configuration

IP Address	IP Mask	Secret
10.10.15.0	255.255.255.0	secret
10.10.8.15	255.255.255.255	newsecret
192.168.25.9	255.255.255.255	mysecret
192.168.102.5	255.255.255.255	999Y2K

Udp Port: 1812

RADIUS Relay Accounting Server Configuration

IP Address	IP Mask	Secret
10.10.1.0	255.255.255.0	N08pxq
192.168.102.5	255.255.255.255	12BE\$56

Udp Port: 1813

**Related
Documentation**

- [Setting a Baseline for RADIUS Relay Statistics on page 301](#)
- [Monitoring RADIUS Relay Server Statistics on page 301](#)
- [Monitoring the Configuration of the RADIUS Relay Server on page 303](#)
- [Monitoring the Status of RADIUS Relay UDP Checksums on page 303](#)

CHAPTER 7

RADIUS Attribute Descriptions

This chapter lists the RADIUS attributes that are supported by JunosE Software. [Table 58 on page 241](#) describes the supported RADIUS IETF attributes. [Table 59 on page 248](#) describes the supported Juniper Networks vendor-specific attributes (VSAs). [Table 60 on page 261](#) describes the DSL Forum VSA formats supported by JunosE Software. [Table 61 on page 262](#) describes RADIUS attributes that are simply passed to their destination by the router.

RADIUS attributes are discussed in the following sections:

- [RADIUS IETF Attributes on page 241](#)
- [Juniper Networks VSAs on page 248](#)
- [DSL Forum VSAs on page 260](#)
- [Pass Through RADIUS Attributes on page 262](#)
- [RADIUS Attributes References on page 262](#)

RADIUS IETF Attributes

[Table 58 on page 241](#) describes the RADIUS IETF attributes supported by JunosE Software. The attributes are sorted by standard number.

Table 58: RADIUS IETF Attributes Supported by JunosE Software

Attribute Number	Attribute Name	Description
[1]	User-Name	<ul style="list-style-type: none">• Name of user to be authenticated• Configurable username override
[2]	User-Password	<ul style="list-style-type: none">• Password of user to be authenticated• Configurable password override• Password Authentication Protocol (PAP)
[3]	CHAP-Password	Response value provided by a Point-to-Point Protocol (PPP) Challenge Handshake Authorization Protocol (CHAP) user in the response to an access challenge

Table 58: RADIUS IETF Attributes Supported by JunosE Software (*continued*)

Attribute Number	Attribute Name	Description
[4]	NAS-IP-Address	<ul style="list-style-type: none"> IP address of the network access server (NAS) that is requesting authentication of the user You can use the <i>radius update-source-addr</i> command to override this behavior.
[5]	NAS-Port	<ul style="list-style-type: none"> Physical port number of the NAS that is authenticating the user See the <i>radius nas-port-format</i>, <i>radius pppoe nas-port-format unique</i>, and <i>radius vlan nas-port-format stacked</i> commands.
[6]	Service-Type	<ul style="list-style-type: none"> Type of service the user has requested or the type of service to be provided Admin, Login, NAS Prompt, or Framed only
[7]	Framed-Protocol	<ul style="list-style-type: none"> Framing protocol used for framed access Standard value of 1 set for PPP Nonstandard value of 1008 set for dynamic ATM
[8]	Framed-IP-Address	<ul style="list-style-type: none"> IP address to be configured for the user 0.0.0.0 or absence is interpreted as 255.255.255.254 See the <i>framed-ip-add acct-start</i> attribute name in the <i>radius include</i> command.
[9]	Framed-IP-Netmask	<ul style="list-style-type: none"> IP network to be configured for the user when the user is a router to a network Absence implies 255.255.255.255
[11]	Filter-Id	<ul style="list-style-type: none"> Name of the filter list for the user Interpreted as input policy name
[12]	Framed-MTU	<ul style="list-style-type: none"> The maximum transmission unit to be configured for the user, when it is not negotiated by some other means (such as PPP). When sent in an Access-Request with an EAP-Message, indicates the maximum size of the EAP-Message string that the external server supports.
[13]	Framed-Compression	Always set to none.
[18]	Reply-Message	<ul style="list-style-type: none"> Text that may be displayed to the user Only the first instance of this attribute is used
[22]	Framed-Route	<p>String that provides routing information to be configured for the user on the NAS; in the format:</p> <pre><addr>[/<maskLen>] <nexthop> [<cost>] [tag <tagValue>] [distance <distValue>]</pre>
[24]	State	<ul style="list-style-type: none"> An arbitrary value that the router includes in new Access-Request packets from the previous Accept-Challenge Applicable for CLI, telnet, or EAP message exchange

Table 58: RADIUS IETF Attributes Supported by JunosE Software (*continued*)

Attribute Number	Attribute Name	Description
[25]	Class	An arbitrary value that the NAS includes in all accounting packets for the user if supplied by the RADIUS server
[26]	Vendor-Specific	Juniper Networks Enterprise number 0x0000130A
[27]	Session-Timeout	Maximum number of consecutive seconds of service to be provided to the user before termination of the session
[28]	Idle-Timeout	Maximum number of consecutive seconds of idle connection provided to the user before termination of the session
[30]	Called-Station-Id	<ul style="list-style-type: none"> Allows the NAS to send the phone number that the user called Not supported for nontunneled or LAC session side For the LNS, the format is the string passed in the Called Number AVP For RADIUS relay server, indicates the subscriber's wireless access point
[31]	Calling-Station-Id	<ul style="list-style-type: none"> Allows the NAS to send the phone number from which the call originated See the <i>radius calling-station-format</i> and the <i>radius calling-station-delimiter</i> commands. For RADIUS relay server, indicates the subscriber's MAC address
[32]	NAS-Identifier	<ul style="list-style-type: none"> Identifies the NAS originating the request System-wide configurable hostname or VR-sensitive configurable NAS-identifier name
[33]	Proxy-State	E Series router's port ID and IP address
[40]	Acct-Status-Type	Indicates whether this Accounting-Request marks the beginning of the user service (Start), the end (Stop), or the interim (Interim-Update)
[41]	Acct-Delay-Time	Indicates how many seconds the client has been trying to send a particular record
[42]	Acct-Input-Octets	<ul style="list-style-type: none"> Indicates how many octets have been received from the port during the time this service has been provided IP subscriber manager—Statistics are reported PPP—Statistics are counted according to the rules of the generic interface MIB
[43]	Acct-Output-Octets	<ul style="list-style-type: none"> Indicates how many octets have been sent to the port during the time this service has been provided IP subscriber manager—Statistics are reported PPP—Statistics are counted according to the rules of the generic interface MIB

Table 58: RADIUS IETF Attributes Supported by JunosE Software (*continued*)

Attribute Number	Attribute Name	Description
[44]	Acct-Session-Id	<ul style="list-style-type: none"> Unique accounting identifier that makes it easy to match start and stop records in a log file See the <i>radius acct-session-id-format</i> and the <i>radius include acct-session-id access-request</i> commands.
[45]	Acct-Authentic	<ul style="list-style-type: none"> Indicates how the user was authenticated: whether by RADIUS, the NAS itself, or another remote authentication protocol Always 1
[46]	Acct-Session-Time	Indicates how long in seconds that the user has received service
[47]	Acct-Input-Packets	<ul style="list-style-type: none"> Indicates how many packets have been received from the port during the time this service has been provided to a framed user IP subscriber manager—Statistics are reported PPP—Statistics are counted according to the rules of the generic interface MIB
[48]	Acct-Output-Packets	<ul style="list-style-type: none"> Indicates how many packets have been sent to the port in the course of delivering this service to a framed user IP subscriber manager—Statistics are reported PPP—Statistics are counted according to the rules of the generic interface MIB
[49]	Acct-Terminate-Cause	<p>Contains the reason the service (a PPP session) was terminated. The service can be terminated for the following reasons:</p> <ul style="list-style-type: none"> User Request (1)—User initiated the disconnect (log out) Idle Timeout (4)—Idle timer has expired Session Timeout (5)—Client reached the maximum continuous time allowed on the service or session Admin Reset (6)—System administrator terminated the session Port Error (8)—PVC failed; no hardware or no interface NAS Error (9)—Negotiation failures, connection failures, or address lease expiration NAS Request (10)—PPP challenge timeout, PPP request timeout, tunnel establishment failure, PPP bundle failure, IP address lease expiration, PPP keep-alive failure, Tunnel disconnect, or an unaccounted-for error
[50]	Acct-Multi-Session-Id	<ul style="list-style-type: none"> String constructed from the Acct-Session-ID of the first PPP link established for the Multilink PPP bundle and the internal Multilink PPP bundle ID. This string is the hexadecimal ASCII characters for two 4-octet unsigned integers. Example: 0a34331200001249.
[51]	Acct-Link-Count	A value that increments with each link that joins the MLPPP bundle. This attribute does not indicate the number of active links. For more details, see RFC 2866—RADIUS Accounting (June 2000).

Table 58: RADIUS IETF Attributes Supported by JunosE Software (*continued*)

Attribute Number	Attribute Name	Description
[52]	Acct-Input-Gigawords	<ul style="list-style-type: none"> Indicates how many times the Acct-Input-Octets counter has wrapped around 2^{32} during the time this service has been provided, and can be present in Accounting-Request records only where the Acct-Status-Type is set to Stop or Interim-Update IP subscriber manager—Statistics are reported PPP—Statistics are counted according to the rules of the generic interface MIB
[53]	Acct-Output-Gigawords	<ul style="list-style-type: none"> Indicates how many times the Acct-Output-Octets counter has wrapped around 2^{32} in the course of delivering this service, and can be present in Accounting-Request records only where the Acct-Status-Type is set to Stop or Interim-Update IP subscriber manager—Statistics are reported PPP—Statistics are counted according to the rules of the generic interface MIB
[55]	Event-Timestamp	Records the time that this event occurred on the NAS, in seconds, since January 1, 1970 00:00 UTC
[60]	CHAP-Challenge	Contains the CHAP challenge sent by the NAS to a PPP CHAP user
[61]	NAS-Port-Type	<ul style="list-style-type: none"> Indicates the type of physical port the NAS is using to authenticate the user See the <i>radius dsl-port-type</i> and the <i>radius ethernet-port-type</i> commands.
[62]	Port-Limit	Specifies the maximum number of MLPPP member links allowed for the subscriber
[64]	Tunnel-Type	<ul style="list-style-type: none"> Which tunneling protocol to use (in the case of a tunnel initiator) or the tunneling protocol in use (in the case of a tunnel terminator) Only L2TP tunnels supported at this time
[65]	Tunnel-Medium-Type	<ul style="list-style-type: none"> Transport medium to use when creating a tunnel for those protocols (such as L2TP) that can operate over multiple transports Only IPv4 supported at this time
[66]	Tunnel-Client-Endpoint	Address of the initiator end of the tunnel
[67]	Tunnel-Server-Endpoint	Address of the server end of the tunnel
[68]	Acct-Tunnel-Connection	<ul style="list-style-type: none"> Indicates the identifier assigned to the tunnel session Value is L2TP call-serial number
[69]	Tunnel-Password	Password to be used to authenticate to a remote server
[77]	Connect-Info	Sent from the NAS to indicate the nature of the user's connection
[79]	EAP-Message	Encapsulates EAP packets, which allows the NAS to authenticate users through EAP without having to understand the EAP protocol

Table 58: RADIUS IETF Attributes Supported by JunosE Software (*continued*)

Attribute Number	Attribute Name	Description
[80]	Message-Authenticator	Must be used in any Access-Request, Access-Accept, Access-Reject or Access-Challenge messages that include EAP-Message attributes
[82]	Tunnel-Assignment-Id	Indicates to the tunnel initiator the particular tunnel to which a session is to be assigned
[83]	Tunnel-Preference	<ul style="list-style-type: none"> If more than one set of tunneling attributes is returned by the RADIUS server to the tunnel initiator, this attribute is included in each set to indicate the relative preference assigned to each tunnel. Included in the Tunnel-Link-Start, the Tunnel-Link-Reject, and the Tunnel-Link-Stop packets (LAC only)
[85]	Acct-Interim-Interval	Number of seconds between each interim accounting update for this session
[86]	Acct-Tunnel-Packets-Lost	Number of packets lost on a given link
[87]	NAS-Port-Id	<ul style="list-style-type: none"> Text string that identifies the physical interface of the NAS that is authenticating the user If the PPP user connects via ATM slot 12, port 2, subinterface 3, vpi 100, vci 101, then the NAS-Port-Id value in the RADIUS packets will be atm 12/2.3:100.101 If the user is a PPP user that started as a result of the E Series LNS feature (that is, no physical port), then the NAS-Port-Id value is as follows: <i>media:local address:peer address:local tunnel id:peer tunnel id:local session id:peer session id:call serial number</i> <ul style="list-style-type: none"> For example: ip:172.81.1.98:172.81.1.99:18d:cb8:ce6:9f4:6 In this case, the local information refers to the LNS, and the peer information refers to the LAC NAS-Port-Id usually contains one of the following: <ul style="list-style-type: none"> atm <slot> / <port><.subinterface>:<vpi>.<vci> FastEthernet <slot> / <port><.subinterface> [<vlan>] GigabitEthernet <slot> / <port><.subinterface> [<vlan>] serial <slot>/<port> [:<sonetPath> [/<sonetTributary (x/x/x)> [/<fractionalInterface>]]] from LNS—ip:local ip:peer ip:local tid:peer tid:local sid:peer sid:call serial number tid—tunnel id sid—session id <p>NOTE: Releases before 4.0.0 did not pass the subinterface number to RADIUS for inclusion in the NAS-Port-Id. If you do not want the subinterface number to be included, you must enter the aaa intf-desc-format include sub-intf disable command to omit the subinterface.</p>
[88]	Framed-Pool	Name of an assigned address pool that should be used to assign an address for the user
[90]	Tunnel-Client-Auth-Id	Name used by the tunnel initiator during the authentication phase of tunnel establishment

Table 58: RADIUS IETF Attributes Supported by JunosE Software (*continued*)

Attribute Number	Attribute Name	Description
[91]	Tunnel-Server-Auth-Id	Name used by the tunnel terminator during the authentication phase of tunnel establishment
[96]	Framed-Interface-Id	IPv6 interface identifier configured by the user
[97]	Framed-Ipv6-Prefix	Provides the IPv6 prefix that is delegated to a downstream CPE
[99]	Framed-Ipv6-Route	Provides routing information to be configured for the user on the NAS
[100]	Framed-Ipv6-Pool	Name of the local address pool from which an IPv6 prefix is assigned to the requesting router
[101]	Error-Cause	4-octet field that contains an integer that specifies the cause of the error
[123]	Delegated-Ipv6-Prefix	IPv6 prefix to be delegated to clients using the DHCPv6 Prefix Delegation mechanism
[135]	Ascend-Primary-DNS	<ul style="list-style-type: none"> Indicates the IP address of the primary DNS The format is 1 byte of type (135), 1 byte of length (length=6), 4 bytes of value (IPv4 address)
[136]	Ascend-Secondary-DNS	<ul style="list-style-type: none"> Indicates the IP address of the secondary DNS The format is 1 byte of type (136), 1 byte of length (length=6), 4 bytes of value (IPv4 address)
[144]	DS-Lite-Tunnel-Name	Specifies the fully qualified domain name (FQDN) of the Address Family Transition Router (AFTR) name to which DHCPv6 client can establish an IPv4-over-IPv6 tunnel (IPv4-over-IPv6 tunnel is commonly referred to as Softwire)
[188]	Ascend-Num-In-Multilink	Current number of links in a multilink bundle
[242]	Ascend-Data-Filter	RADIUS policy definitions used to configure a policy to classify packet flows and perform filter, forward, packet marking, rate-limit profile, and traffic class actions

- Related Documentation**
- [RADIUS IETF Attributes Supported for Subscriber AAA Access Messages on page 183](#)
 - [RADIUS IETF Attributes Supported for Subscriber AAA Accounting Messages on page 193](#)
 - [RADIUS IETF Attributes Supported for AAA Tunnel Accounting Messages on page 200](#)
 - [CLI Commands Used to Configure RADIUS IETF Attributes on page 209](#)

Juniper Networks VSAs

Table 59 on page 248 lists Juniper Networks VSA formats for RADIUS. JunosE Software uses the vendor ID assigned to Juniper Networks (vendor ID 4874) by the Internet Assigned Numbers Authority (IANA).

Table 59: Juniper Networks (Vendor ID 4874) VSA Formats

Attribute Number	Attribute Name	Description	Length	Subtype Length	Value
[26-1]	Virtual-Router	<ul style="list-style-type: none"> Virtual router name for the Broadband Remote Access Server (B-RAS) user's IP interface. Allowed only from RADIUS server in default virtual router context. For restricted users, specifies the only virtual router that the user can access. For nonrestricted users, specifies the initial virtual router that the user accesses. For tunneled connections, specifies the tunnel source parameter where the source address for the tunneled connection is resolved. See the enable command in the <i>Passwords and Security</i> chapter in the <i>JunosE System Basics Configuration Guide</i>. 	len	sublen	string: virtual-router-name
[26-2]	Local-Address-Pool	<ul style="list-style-type: none"> Name of an assigned address pool that should be used to assign an address for the user Same as RADIUS attribute 88, Framed-Pool 	len	sublen	string: address-pool-name
[26-3]	Local-Interface	<p>Interface to apply to the E Series side of the connection</p> <p>The interface value can be one of the following:</p> <ul style="list-style-type: none"> The IP address (with subnet mask) The loopback interface 	len	sublen	string: local-interface
[26-4]	Primary-DNS	<ul style="list-style-type: none"> B-RAS user's DNS address negotiated during IPCP 4-octet IP address 	12	6	integer: 4-byte primary-dns-address
[26-5]	Secondary-DNS	<ul style="list-style-type: none"> B-RAS user's DNS address negotiated during IPCP 4-octet IP address 	12	6	integer: 4-byte secondary-dns-address

Table 59: Juniper Networks (Vendor ID 4874) VSA Formats (*continued*)

Attribute Number	Attribute Name	Description	Length	Subtype Length	Value
[26-6]	Primary-WINS (NBNS)	<ul style="list-style-type: none"> B-RAS user's WINS (NBNS) address negotiated during IPCP 4-octet IP address 	12	6	integer: 4-byte primary-wins-address
[26-7]	Secondary-WINS (NBNS)	<ul style="list-style-type: none"> B-RAS user's WINS (NBNS) address negotiated during IPCP 4-octet IP address 	12	6	integer: 4-byte secondary-wins-address
[26-8]	Tunnel-Virtual-Router	For tunneled connections, specifies the virtual router associated with the tunnel connection	len	sublen	string: tunnel-virtual-router
[26-9]	Tunnel-Password	Tunnel password in cleartext	len	sublen	string: tunnel-password
[26-10]	Ingress-Policy-Name	IPv4 input policy name to apply to B-RAS user's interface	len	sublen	string: input-policy-name
[26-11]	Egress-Policy-Name	IPv4 output policy name to apply to B-RAS user's interface	len	sublen	string: output-policy-name
[26-12]	Ingress-Statistics	Enable or disable input statistics on B-RAS user's interface	12	6	integer: 0 = disable, 1 = enable
[26-13]	Egress-Statistics	Enable or disable output statistics on B-RAS user's interface	12	6	integer: 0 = disable, 1 = enable
[26-14]	Service-Category	ATM service category to apply to B-RAS user's interface	12	6	integer: 1= UBR, 2 = UBR PCR, 3 = NRT VBR, 4 = CBR 5 = RT VBR,
[26-15]	PCR	<ul style="list-style-type: none"> Peak cell rate 4-octet integer 	12	6	integer: 4-octet
[26-16]	SCR	<ul style="list-style-type: none"> Sustained cell rate 4-octet integer 	12	6	integer: 4-octet
[26-17]	Mbs	<ul style="list-style-type: none"> Maximum burst rate 4-octet integer 	12	6	integer: 4-octet
[26-18]	Init-CLI-Access-Level	<ul style="list-style-type: none"> Specifies the initial level of access to CLI commands See the enable command in the <i>Passwords and Security</i> chapter in the <i>JunosE System Basics Configuration Guide</i>. 	len	sublen	single attribute: enter 0, 1, 5, 10, or 15

Table 59: Juniper Networks (Vendor ID 4874) VSA Formats (*continued*)

Attribute Number	Attribute Name	Description	Length	Subtype Length	Value
[26-19]	Allow-All-VR-Access	<ul style="list-style-type: none"> Specifies user access to all virtual routers See the enable command in the <i>Passwords and Security</i> chapter in the <i>JunosE System Basics Configuration Guide</i>. 	len	sublen	integer: 0 = disable, 1 = enable
[26-20]	Alt-CLI-Access-Level	<ul style="list-style-type: none"> Specifies other levels of access to CLI commands See the enable command in the <i>Passwords and Security</i> chapter of the <i>JunosE System Basics Configuration Guide</i>. 	len	sublen	single attribute; enter 0, 1, 5, 10, or 15
[26-21]	Alt-CLI-Vrouter-Name	<ul style="list-style-type: none"> For restricted users, specifies other VRs that the user may access. See the enable command in the <i>Passwords and Security</i> chapter in the <i>JunosE System Basics Configuration Guide</i>. 	len	sublen	string: virtual-router-name
[26-22]	Sa-Validate	<ul style="list-style-type: none"> Enable or disable source address validation on a user's interface 4-octet integer 	len	sublen	integer: 0 = disable, 1 = enable
[26-23]	Igmp-Enable	<ul style="list-style-type: none"> Enable or disable IGMP on a user's interface Allows the end user to register for the reception of multicast services 4-octet integer 	len	sublen	integer: 0 = disable, 1 = enable
[26-24]	Pppoe-Description	The string <i>pppoe <mac addr></i> sent to the RADIUS server supplied by PPPoE	len	sublen	string: pppoe<mac addr>
[26-25]	Redirect-Vrouter-Name	<ul style="list-style-type: none"> Virtual router name indicating the VR context in which to authenticate the user Behavior is similar to that of a remote domain-map lookup. 	len	sublen	authentication-redirection
[26-26]	QoS-Profile-Name	Name of the QoS profile to attach to the user's interface	len	sublen	string: qos-profile-name
[26-28]	Pppoe-Url	PPPoE URL that is passed to PPPoE subscribers	len	sublen	string:URL
[26-30]	Tunnel-Nas-Port-Method	Conveys nasPort and nasPort type in tunnel	12	6	4-octet integer: 0 = none, 1 = Cisco CLID
[26-31]	Service-Bundle	Specifies the SRC service bundle	len	sublen	string

Table 59: Juniper Networks (Vendor ID 4874) VSA Formats (*continued*)

Attribute Number	Attribute Name	Description	Length	Subtype Length	Value
[26-33]	Tunnel-Max-Sessions	Maximum number of sessions allowed in a tunnel	12	6	integer: 4-octet
[26-34]	Framed-Ip-Route-Tag	Route tag to apply to returned framed-ip-address	12	6	integer: 4-octet
[26-35]	Tunnel-Dialout-Number	Dial number in L2TP dial-out	len	sublen	string:dial-out-number
[26-36]	PPP-Username	Username used in PPP L2TP dial-out sessions at the LNS for L2TP dial-out	len	sublen	string: ppp-username
[26-37]	PPP-Password	Password used in PPP L2TP dial-out sessions at the LNS for L2TP dial-out	len	sublen	string: ppp-password
[26-38]	PPP-Protocol	PPP authentication protocol used for L2TP dial-out sessions at the LNS	12	6	integer: 0 = none; 1 = PAP; 2 = CHAP; 3 = PAP-CHAP; 4 = CHAP-PAP
[26-39]	Tunnel-Min-Bps	Minimum line speed for L2TP dial-out	12	6	integer
[26-40]	Tunnel-Max-Bps	Maximum line speed for L2TP dial-out	12	6	integer
[26-41]	Tunnel-Bearer-Type	Bearer capability required for L2TP dial-out	12	6	integer: 0 = none; 1= analog; 2 = digital
[26-42]	Input-GigaPkts	Number of times input-packets attribute rolls over its 4-octet field	12	6	integer
[26-43]	Output-GigaPkts	Number of times output-packets attribute rolls over its 4-octet field	12	6	integer
[26-44]	Tunnel-Interface-Id	Tunnel interface selector that AAA caches as part of the tunnel-session profile and the user's profile. This attribute is available to the RADIUS authentication and accounting servers.	len	sublen	string: tunnel selector
[26-45]	Ipv6-Virtual-Router	Virtual router name for B-RAS user's IPv6 interface	len	sublen	string: virtual-router-name
[26-46]	Ipv6-Local-Interface	Local IPv6 interface to apply to the E Series side of the connection	len	sublen	string: ipv6-local-interface
[26-47]	Ipv6-Primary-DNS	B-RAS user's primary IPv6 DNS address negotiated by DHCP	len	sublen	hexadecimal string: ipv6-primary-dns-address

Table 59: Juniper Networks (Vendor ID 4874) VSA Formats (*continued*)

Attribute Number	Attribute Name	Description	Length	Subtype Length	Value
[26-48]	Ipv6-Secondary-DNS	B-RAS user's secondary IPv6 DNS address negotiated by DHCP	len	sublen	hexadecimal string: ipv6-primary-dns-address
[26-51]	Disconnect-Cause	L2TP PPP disconnect cause information received by the LAC	len	sublen	string:l2tp-ppp-disconnect-cause
[26-52]	Radius-Client-Address	RADIUS relay server's IP address	12	6	integer:4-octet
[26-53]	Service-Description	AAA profile service description string	len	sublen	string:profile-service-description
[26-54]	L2tp-Recv-Window-Size	<ul style="list-style-type: none"> L2TP receive window size (RWS) for a tunnel on the LAC Number of packets that the peer can transmit without receiving an acknowledgment from the router 4-octet integer 	12	6	integer:4-octet
[26-55]	DHCP-Options	Client's DHCP options	len	sublen	string:dhcp-options
[26-56]	DHCP-MAC-Address	Client's MAC address	len	sublen	string:mac-address
[26-57]	DHCP-GI-Address	DHCP relay agent's IP address	12	6	integer:4-octet
[26-58]	LI-Action	Packet mirroring action	len	sublen	Salt encrypted integer: 0 = stop monitoring; 1 = start monitoring; 2 = no action
[26-59]	Med-Dev-Handle	Hexadecimal string used to determine mirror header attributes, prepended to each mirrored packet that is sent to the analyzer device	len	sublen	Salt encrypted string; hexadecimal string of 4 bytes or 8 bytes
[26-60]	Med-Ip-Address	IP address of analyzer device to which mirrored packets are forwarded	len	sublen	Salt encrypted IP address
[26-61]	Med-Port-Number	UDP port in the analyzer device to which mirrored packets are forwarded	len	sublen	Salt encrypted integer
[26-62]	MLPPP-Bundle-Name	Text string that identifies the Multilink PPP bundle name	len	sublen	string:mlppp-bundle-name
[26-63]	Interface-Desc	Text string that identifies the subscriber's access interface	len	sublen	string:interface-description

Table 59: Juniper Networks (Vendor ID 4874) VSA Formats (*continued*)

Attribute Number	Attribute Name	Description	Length	Subtype Length	Value
[26-64]	Tunnel-Group	Name of the tunnel group assigned to a domain map	len	sublen	string:tunnel-group-name
[26-65]	Activate-Service	Service to activate for the subscriber	len	sublen	string:service-name
[26-66]	Deactivate-Service	Service to deactivate for the subscriber	len	sublen	string:service-name
[26-67]	Service-Volume-tagX	Amount of traffic, in MB, that can use the service; service is deactivated when the volume is exceeded	12	6	integer: volume in MB; 0 = infinite volume
[26-68]	Service-Timeout-tagX	Number of seconds that the service can be active; service is deactivated when the timeout expires	12	6	integer: time in seconds; 0 = no timeout
[26-69]	Service-Statistics-tagX	Enable or disable statistics for the service	12	6	integer: 0 = disable; 1 = enable time statistics; 2 = enable time and volume statistics
[26-70]	Ignore-DF-Bit	Enable or disable the ignore don't fragment (DF) bit feature on a B-RAS user's interface	12	6	integer: 0 = disable; 1 = enable
[26-71]	IGMP-Access-Name	Access List to use for the group (G) filter	len	sublen	string:32-octet
[26-72]	IGMP-Access-Src-Name	Access List to use for the source-group (S,G) filter	len	sublen	string:32-octet
[26-73]	IGMP-OIF-Map-Name	Multicast OIF (outgoing interface) mapping	len	sublen	string:32-octet
[26-74]	MLD-Access-Name	Access List to use for the group (G) filter	len	sublen	string:32-octet
[26-75]	MLD-Access-Src-Name	Access List to use for the source-group (S,G) filter	len	sublen	string:32-octet
[26-76]	MLD-OIF-Map-Name	Multicast OIF (outgoing interface) mapping	len	sublen	string:32-octet
[26-77]	MLD-Version	MLD Protocol Version (MLD Version 1 = 1; MLD Version 2 = 2)	12	6	integer:1-octet
[26-78]	IGMP-Version	IGMP Protocol Version (IGMP Version 1=1; IGMP Version 2 = 2; IGMP Version 3 = 3)	12	6	integer:1-octet
[26-79]	IP-Mcast-Adm-Bw-Limit	The maximum multicast bandwidth that will be admitted on an IP interface, in Kbps	12	6	integer:4-octet

Table 59: Juniper Networks (Vendor ID 4874) VSA Formats (*continued*)

Attribute Number	Attribute Name	Description	Length	Subtype Length	Value
[26-80]	IPv6-Mcast-Adm-Bw-Limit	The maximum multicast bandwidth that will be admitted on an IPv6 interface, in Kbps	12	6	integer:4-octet
[26-81]	L2c-Information	Series of type length value (tlv) fields (binary) representing the access loop parameters as defined in GSMP extensions for layer2 control (L2C) Topology Discovery and Line Configuration—draft-wadhwa-gsmp-l2control-configuration-00.txt (July 2006 expiration)	len	sublen	string: format is a series of type length value (tlv) fields (binary) representing the access loop parameters
[26-82]	Qos-Parameters	Name of the QoS parameter instance to create on the user's interface, followed by the value of the parameter. For example, the max-bandwidth 4000000 parameter instance represents the parameter name that was defined using the qos-parameter-define command (max-bandwidth) and the value to assign to the parameter (4000000). Multiple instances of this VSA can be returned from RADIUS using this format.	len	sublen	string: format is <i>parameter name parameter value</i> , where <i>parameter name</i> is ASCII name of a parameter name found in the QoS parameter definition and <i>parameter value</i> is the ASCII representation of 0–21474836470; multiple instances of this VSA can be returned from RADIUS using this format
[26-83]	Service-Session	Name of the service (including parameter values) that is associated with service manager statistics	len	sublen	string:service-name
[26-84]	Mobile-IP-Algorithm	Authentication algorithm used for Mobile IP registration	12	6	integer: 4-octet
[26-85]	Mobile-IP-SPI	Security parameter index for Mobile IP registration	12	6	integer: 4-octet
[26-86]	Mobile-IP-Key	Security association MD-5 key for Mobile IP registration	len	sublen	string: 32-octet
[26-87]	Mobile-IP-Replay	Replay time stamp for Mobile IP registration	12	6	integer: 4-octet
[26-88]	Mobile-IP-Access-Control-List	Access control list to filter on basis of care-of address	len	sublen	string: 32-octet
[26-89]	Mobile-IP-Lifetime	Registration lifetime for Mobile IP registration	12	6	integer: 4-octet

Table 59: Juniper Networks (Vendor ID 4874) VSA Formats (*continued*)

Attribute Number	Attribute Name	Description	Length	Subtype Length	Value
[26-90]	L2TP-Resynch-Method	L2TP peer resynchronization method	12	6	integer: 0 = disabled; 1 = failover protocol; 2 = silent failover; 3 = failover protocol with silent failover as backup
[26-91]	Tunnel-Switch-Profile	<ul style="list-style-type: none"> Name of the L2TP tunnel switch profile The L2TP tunnel switch profile defines the L2TP tunnel switching behavior for the interfaces to which this profile is assigned 	len	sublen	string: tunnel-switch-profile
[26-92]	L2C-Up-Stream-Data	Actual upstream rate access loop parameter (ASCII encoded) as defined in GSMP extensions for layer2 control (L2C) Topology Discovery and Line Configuration—draft-wadhwa-gsmp-l2control-configuration-00.txt (July 2006 expiration).	len	sublen	string: actual upstream rate access loop parameter (ASCII encoded)
[26-93]	L2C-Down-Stream-Data	Actual downstream rate access loop parameter (ASCII encoded) as defined in GSMP extensions for layer2 control (L2C) Topology Discovery and Line Configuration—draft-wadhwa-gsmp-l2control-configuration-00.txt (July 2006 expiration).	len	sublen	string: actual downstream rate access loop parameter (ASCII encoded)
[26-94]	Tunnel-Tx-Speed-Method	The method that the router uses to calculate the transmit connect speed of the subscriber's access interface. This speed is reported in L2TP Transmit (TX) Speed AVP 24. During the establishment of an L2TP tunnel session, the LAC sends AVP 24 to the LNS to convey the transmit speed of the subscriber's access interface.	12	6	integer: 1 = static-layer2, TX speed based on static layer 2 settings; 2 = dynamic-layer2, TX speed based on dynamic layer 2 settings; 3 = qos, TX speed based on QoS settings; 4 = actual, TX speed that is the lesser of the dynamic-layer2 value or the qos value
[26-95]	IGMP-Query-Interval	IGMP Query Interval	12	6	integer: 4-octet
[26-96]	IGMP-Max-Resp-Time	IGMP Maximum Response Time	12	6	integer: 4-octet
[26-97]	IGMP-Immediate-Leave	IGMP Immediate Leave	12	6	4-octet integer: 0 = disabled 1 = enabled

Table 59: Juniper Networks (Vendor ID 4874) VSA Formats (continued)

Attribute Number	Attribute Name	Description	Length	Subtype Length	Value
[26-98]	MLD-Query-Interval	MLD Query Interval	12	6	integer: 4-octet
[26-99]	MLD-Max-Resp-Time	MLD Maximum Response Time	12	6	integer: 4-octet
[26-100]	MLD-Immediate-Leave	MLD Immediate Leave	12	6	integer: 4-octet; 0 = disabled 1 = enabled
[26-101]	IP-Block-Multicast	Block all multicast traffic with a scope larger than link-local (for example, global) and prevent mroute creation under these conditions. This attribute does not affect reception of link-local multicast packets.	12	6	integer: 4-octet; 0 = disabled; 1 = enabled
[26-102]	IGMP-Explicit-Tracking	Enable or disable explicit host tracking for IPv4 IGMP interfaces. This option enables the router to explicitly track each individual host that is joined to a group or channel on a particular multi-access network.	12	6	integer: 4-octet; 0 = disabled; 1 = enabled
[26-103]	IGMP-No-Tracking-V2-Grps	Disable IGMP explicit host tracking for groups that contain IGMP V2 hosts. This attribute is valid only if IGMP V3 is enabled on the interface.	12	6	integer: 4-octet; 0 = disabled; 1 = enabled
[26-104]	MLD-Explicit-Tracking	Enable or disable explicit host tracking for IPv6 MLD interfaces. This option enables the router to explicitly track each individual host that is joined to a group or channel on a particular multi-access network.	12	6	integer: 4-octet; 0 = disabled; 1 = enabled
[26-105]	MLD-No-Tracking-V1-Grps	Disable MLD explicit host tracking for groups that contain MLD V1 hosts. This attribute is valid only if MLD V2 is enabled on the interface.	12	6	integer: 4-octet; 0 = disabled; 1 = enabled
[26-106]	Ipv6-Ingress-Policy-Name	IPv6 ingress policy that is applied to the subscriber interface	len	sublen	string: Ipv6-Ingress-Policy-Name
[26-107]	Ipv6-Egress-Policy-Name	IPv6 egress policy that is applied to the subscriber interface	len	sublen	string: Ipv6-Egress-Policy-Name
[26-110]	Acc-Loop-Cir-Id	Identification of the subscriber node connection to the access node	len	sublen	string: up to 63 ASCII characters
[26-111]	Acc-Aggr-Cir-Id-Bin	Unique identification of the DSL line	len	sublen	integer: 8-octet

Table 59: Juniper Networks (Vendor ID 4874) VSA Formats (*continued*)

Attribute Number	Attribute Name	Description	Length	Subtype Length	Value
[26-112]	Acc-Aggr-Cir-Id-Asc	Identification of the uplink on the access node. For example: <ul style="list-style-type: none"> For Ethernet access aggregation: <i>ethernet slot/port [:inner-vlan-id] [:outer-vlan-id]</i> For ATM aggregation: <i>atm slot/port:vpi.vci</i> 	len	sublen	string: up to 63 ASCII characters
[26-113]	Act-Data-Rate-Up	Actual upstream data rate of the subscriber's synchronized DSL link	12	6	integer: 4-octet
[26-114]	Act-Data-Rate-Dn	Actual downstream data rate of the subscriber's synchronized DSL link	12	6	integer: 4-octet
[26-115]	Min-Data-Rate-Up	Minimum upstream data rate configured for the subscriber	12	6	integer: 4-octet
[26-116]	Min-Data-Rate-Dn	Minimum downstream data rate configured for the subscriber	12	6	integer: 4-octet
[26-117]	Att-Data-Rate-Up	Upstream data rate that the subscriber can attain	12	6	integer: 4-octet
[26-118]	Att-Data-Rate-Dn	Downstream data rate that the subscriber can attain	12	6	integer: 4-octet
[26-119]	Max-Data-Rate-Up	Maximum upstream data rate configured for the subscriber	12	6	integer: 4-octet
[26-120]	Max-Data-Rate-Dn	Maximum downstream data rate configured for the subscriber	12	6	integer: 4-octet
[26-121]	Min-LP-Data-Rate-Up	Minimum upstream data rate in low power state configured for the subscriber	12	6	integer: 4-octet
[26-122]	Min-LP-Data-Rate-Dn	Minimum downstream data rate in low power state configured for the subscriber	12	6	integer: 4-octet
[26-123]	Max-Interlv-Delay-Up	Maximum one-way upstream interleaving delay configured for the subscriber	12	6	integer: 4-octet
[26-124]	Act-Interlv-Delay-Up	Subscriber's actual one-way upstream interleaving delay	12	6	integer: 4-octet
[26-125]	Max-Interlv-Delay-Dn	Maximum one-way downstream interleaving delay configured for the subscriber	12	6	integer: 4-octet

Table 59: Juniper Networks (Vendor ID 4874) VSA Formats (*continued*)

Attribute Number	Attribute Name	Description	Length	Subtype Length	Value
[26-126]	Act-Interlv-Delay-Dn	Subscriber's actual one-way downstream interleaving delay	12	6	integer: 4-octet
[26-127]	DSL-Line-State	State of the DSL line	12	6	4-octet integer 1 = Show uptime 2 = Idle 3 = Silent
[26-128]	DSL-Type	Encapsulation used by the subscriber associated with the DSLAM interface from which requests are initiated	11	5	string: 3-byte
[26-129]	Ipv6-NdRa-Prefix	Prefix value in IPv6 Neighbor Discovery route advertisements	len	sublen	hexadecimal string
[26-130]	QoS-Interfaceset-Name	Name of the QoS interface set to attach to the subscriber interface	len	sublen	string: qos-interfaceset-name
[26-140]	Service-Interim-Acct-Interval	Amount of time between interim accounting updates for this service.	12	6	integer: time in the range 600–86400 seconds; 0 = disabled
[26-141]	Downstream-Calculated-Qos-Rate	Calculated downstream QoS rate in Kbps as set by the ANCP configuration	12	6	integer: 4-octet
[26-142]	Upstream-Calculated-Qos-Rate	Calculated downstream QoS rate in Kbps as set by the ANCP configuration	12	6	integer: 4-octet
[26-143]	Max-Clients-Per-Interface	Maximum number of PPPoE client sessions supported per interface. For DHCP clients, this value is the maximum number of PPPoE sessions per logical interface. For PPPoE, this value is the maximum number of PPPoE subinterfaces per a PPPoE major interface. <i>See JunosE Release Notes, Appendix A, System Maximums</i> corresponding to your software release for information about the maximum number of PPPoE subinterfaces supported for each line module.	12	6	integer: 4-octet
[26-144]	PPP-Monitor-Ingress-Only	Enable or disable monitoring of only ingress traffic to determine inactivity of a PPP session and subsequent disconnection of an inactive session. If this option is disabled or not configured, the router monitors both ingress traffic and egress traffic to determine session inactivity.	12	6	integer: 0 = disable, 1 = enable

Table 59: Juniper Networks (Vendor ID 4874) VSA Formats (*continued*)

Attribute Number	Attribute Name	Description	Length	Subtype Length	Value
[26-147]	Backup-Address-Pool	<p>Name of the backup local address pool that can be used to assign addresses to users being authenticated by a RADIUS server, when the existing addresses in the primary local address pool are fully exhausted.</p> <p>The authentication server overrides the backup local address pool name configured using this attribute with the backup local address pool name received in the RADIUS-Access-Accept message.</p>	len	sublen	string: Backup-address-pool-name
[26-150]	ICR-Partition-Id	Used in all the RADIUS authentication and accounting (Acct-Start, Acct-Stop, and Interim-Acct messages for both user and service accounting) messages corresponding to a subscriber to determine the partition in which the subscriber has logged in	len	sublen	string:icr-partition-id
[26-151]	Ipv6-Acct-Input-Octets	Number of times that IPv6 octets have been received from the port during the time this service has been provided	12	6	4-octet integer
[26-152]	Ipv6-Acct-Output-Octets	Number of times that IPv6 octets have been sent to the port during the time this service has been provided	12	6	4-octet integer
[26-153]	Ipv6-Acct-Input-Packets	Number of times that IPv6 packets have been received from the port during the time this service has been provided to a framed user	12	6	4-octet integer
[26-154]	Ipv6-Acct-Output-Packets	Number of times that IPv6 packets have been sent to the port in the course of delivering this service to a framed user	12	6	4-octet integer
[26-155]	Ipv6-Acct-Input-Gigawords	Number of times that the IPv6-Acct-Input-Octets counter has wrapped around 2^{32} during the time this service has been provided, and can be present in Accounting-Request records only where the Acct-Status-Type is set to Stop or Interim-Update	12	6	4-octet integer
[26-156]	Ipv6-Acct-Output-Gigawords	Number of times that the IPv6-Acct-Output-Octets counter has wrapped around 2^{32} in the course of delivering this service, and can be present in Accounting-Request records only where the Acct-Status-Type is set to Stop or Interim-Update	12	6	4-octet integer

Table 59: Juniper Networks (Vendor ID 4874) VSA Formats (*continued*)

Attribute Number	Attribute Name	Description	Length	Subtype Length	Value
[26-157]	Ipv6-Ndra-Pool	Used in RADIUS Access-Accept message to inform the E Series router to allocate IPv6 Neighbor Discovery router advertisement prefix from this pool for the subscriber. If CLI knob aaa dhcpv6-ndra-pool override is disabled, JunosE interprets this attribute as Neighbor Discovery router advertisement local address pool name.	len	sublen	String: 16 alpha-numeric characters
[26-161]	Delegated-Ipv6-Pool	Used in RADIUS Access-Accept message to inform the E Series router to allocate IPv6 Neighbor Discovery router advertisement prefix from this pool for the subscriber. If CLI knob aaa dhcpv6-ndra-pool override is enabled, JunosE interprets this attribute as DHCPV6 PD pool name.	len	sublen	String: 16 alpha-numeric characters
[26-164]	Ipv4-release-control	Causes the PPP application to notify the RADIUS server regarding IPv4 addresses released by a subscriber in a dual-stack network, when an IPCP negotiation for IPv4 sessions is terminated or if the IPv4 session becomes inactive. This attribute is added to RADIUS messages only if the subscriber session is of a dual-stack type and if the IPv4 address is allocated from the RADIUS server and not from local address pools.	len	sublen	String: 32 alpha-numeric characters
[26-165]	PCP-Server-Name	Specifies the PCP server name to which DHCP clients send PCP requests. A PCP client must know the fully qualified domain name (FQDN) of a PCP server, before it can communicate with the latter in order to perform the relevant PCP functions.	len	sublen	String: 245 octets(alpha-numeric characters, dashes, periods)

Related Documentation

- [Juniper Networks VSAs Supported for Subscriber AAA Access Messages on page 186](#)
- [Juniper Networks VSAs Supported for Subscriber AAA Accounting Messages on page 196](#)
- [CLI Commands Used to Configure Juniper Networks VSAs on page 213](#)
- [CLI Commands Used to Include ANCP-Related Juniper Networks VSAs in Access and Accounting Messages on page 215](#)

DSL Forum VSAs

Table 60 on page 261 describes the DSL Forum VSAs supported by JunosE Software for RADIUS. JunosE Software uses the vendor ID assigned to the DSL Forum (3561, or DE9 in hexadecimal format) by the Internet Assigned Numbers Authority (IANA).

Table 60: JunosE Software DSL Forum (Vendor ID 3561) VSA Formats

Attribute Number	Attribute Name	Description	Length	Subtype Length	Value
[26-1]	Agent-Circuit-Id	Identifier for the subscriber agent circuit ID that corresponds to the DSLAM interface from which subscriber requests are initiated	len	sublen	string: agent-circuit-id
[26-2]	Agent-Remote-Id	Unique identifier for the subscriber associated with the DSLAM interface from which requests are initiated	len	sublen	string: agent-remote-id
[26-129]	Actual-Data-Rate-Upstream	Actual upstream data rate of the subscriber's synchronized DSL link	12	6	integer: 4-octet
[26-130]	Actual-Data-Rate-Downstream	Actual downstream data rate of the subscriber's synchronized DSL link	12	6	integer: 4-octet
[26-131]	Minimum-Data-Rate-Upstream	Minimum upstream data rate configured for the subscriber	12	6	integer: 4-octet
[26-132]	Minimum-Data-Rate-Downstream	Minimum downstream data rate configured for the subscriber	12	6	integer: 4-octet
[26-133]	Attainable-Data-Rate-Upstream	Upstream data rate that the subscriber can attain	12	6	integer: 4-octet
[26-134]	Attainable-Data-Rate-Downstream	Downstream data rate that the subscriber can attain	12	6	integer: 4-octet
[26-135]	Maximum-Data-Rate-Upstream	Maximum upstream data rate configured for the subscriber	12	6	integer: 4-octet
[26-136]	Maximum-Data-Rate-Downstream	Maximum downstream data rate configured for the subscriber	12	6	integer: 4-octet
[26-137]	Minimum-Data-Rate-Upstream-Low-Power	Minimum upstream data rate in low power state configured for the subscriber	12	6	integer: 4-octet
[26-138]	Minimum-Data-Rate-Downstream-Low-Power	Minimum downstream data rate in low power state configured for the subscriber	12	6	integer: 4-octet
[26-139]	Maximum-Interleaving-Delay-Upstream	Maximum one-way upstream interleaving delay configured for the subscriber	12	6	integer: 4-octet
[26-140]	Actual-Interleaving-Delay-Upstream	Subscriber's actual one-way upstream interleaving delay	12	6	integer: 4-octet
[26-141]	Maximum-Interleaving-Delay-Downstream	Maximum one-way downstream interleaving delay configured for the subscriber	12	6	integer: 4-octet

Table 60: JunosE Software DSL Forum (Vendor ID 3561) VSA Formats
(continued)

Attribute Number	Attribute Name	Description	Length	Subtype Length	Value
[26-142]	Actual-Interleaving-Delay-Downstream	Subscriber's actual one-way downstream interleaving delay	12	6	integer: 4-octet
[26-144]	Access-Loop-Encapsulation	Encapsulation used by the subscriber associated with the DSLAM interface from which requests are initiated	11	5	string: 3-byte
[26-254]	IWF-Session	Indication that the interworking function (IWF) has been performed for the subscriber's session to enable the transport of PPP over ATM traffic on a PPPoE interface	8	2	No data field required

Related Documentation

- [DSL Forum VSAs Supported for AAA Access and Accounting Messages on page 206](#)
- [DSL Forum VSAs in AAA Access and Accounting Messages Overview on page 206](#)
- [CLI Commands Used to Include DSL Forum VSAs in Access and Accounting Messages on page 217](#)

Pass Through RADIUS Attributes

[Table 61 on page 262](#) describes the RADIUS attribute that is not processed by JunosE Software. The router simply passes this attribute to its destination.

Table 61: RADIUS Attribute Passed Through by JunosE Software

Standard Number	Attribute Name	Description
[79]	EAP-Message	<ul style="list-style-type: none"> • Used by RADIUS relay servers • Passed through to the RADIUS server

Related Documentation

- [RADIUS IETF Attributes Supported for Subscriber AAA Access Messages on page 183](#)

RADIUS Attributes References

For more information about RADIUS attributes, see the following RFCs:

- RFC 2661—Layer Two Tunneling Protocol “L2TP” (August 1999)
- RFC 2865—Remote Authentication Dial In User Service (RADIUS) (June 2000)
- RFC 2866—RADIUS Accounting (June 2000)
- RFC 2867—RADIUS Accounting Modifications for Tunnel Protocol Support (June 2000)

- RFC 2868—RADIUS Attributes for Tunnel Protocol Support (June 2000)
- RFC 2869—RADIUS Extensions (June 2000)
- RFC 3748—Extensible Authentication Protocol (EAP) (June 2004)
- RFC 4679—DSL Forum Vendor-Specific RADIUS Attributes (September 2006)



NOTE: IETF drafts are valid for only 6 months from the date of issuance. They must be considered as works in progress. Please refer to the IETF website at <http://www.ietf.org> for the latest drafts.

**Related
Documentation**

- [RADIUS References on page 181](#)
- [RADIUS IETF Attributes on page 241](#)
- [Juniper Networks VSAs on page 248](#)
- [DSL Forum VSAs on page 260](#)

CHAPTER 8

Application Terminate Reasons

This chapter lists the default mappings for application terminate reasons to RADIUS Acct-Terminate-Cause attributes. [Table 62 on page 265](#) lists the default mappings for AAA, [Table 63 on page 267](#) lists default mappings for L2TP, [Table 64 on page 283](#) lists the default mappings for PPP, and [Table 65 on page 290](#) lists default mappings for RADIUS client. See “[Mapping Application Terminate Reasons and RADIUS Terminate Codes Overview](#)” on [page 36](#) in the *Configuring Remote Access* chapter. For information about configuring custom mappings for application terminate reasons to RADIUS Acct-Terminate-Cause attributes.

- [AAA Terminate Reasons on page 265](#)
- [L2TP Terminate Reasons on page 266](#)
- [PPP Terminate Reasons on page 283](#)
- [RADIUS Client Terminate Reasons on page 290](#)

AAA Terminate Reasons

[Table 62 on page 265](#) lists the default AAA terminate mappings. The table indicates the supported AAA terminate and deny reasons and the RADIUS Acct-Terminate-Cause attributes they are mapped to by default.

Table 62: Default AAA Mappings

AAA Shutdown or Deny Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
deny address allocation failure	17	user error
deny address assignment failure	17	user error
deny application error	17	user error
deny authentication denied	17	user error
deny authentication failure	17	user error
deny authorization failure	17	user error

Table 62: Default AAA Mappings (*continued*)

AAA Shutdown or Deny Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
deny incompatible request	17	user error
deny invalid tunnel configuration	17	user error
deny limit exceeded	17	user error
deny mixed user types	10	nas request
deny no access challenge support	17	user error
deny no address allocation resources	17	user error
deny no resources	10	nas request
deny redirected authentication failure	17	user error
deny server not available	17	user error
deny server request timeout	17	user error
deny terminating user	10	nas request
deny unknown subscriber	17	user error
deny user termination	17	user error
shutdown address lease expiration	10	nas request
shutdown administrative reset	6	admin reset

Related Documentation

- [Mapping Application Terminate Reasons and RADIUS Terminate Codes Overview on page 36](#)
- [Monitoring Application Terminate Reason Mappings on page 167](#)

L2TP Terminate Reasons

Table 63 on page 267 lists the default L2TP terminate mappings. The table indicates the supported L2TP terminate reasons and the RADIUS Acct-Terminate-Cause attributes they are mapped to by default.

Table 63: Default L2TP Mappings

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
session access interface down	8	port error
session admin close	6	admin reset
session admin drain	6	admin reset
session call down	10	nas request
session call failed	15	service unavailable
session create failed limit reached	9	nas error
session create failed no resources	9	nas error
session create failed single shot tunnel already fired	9	nas error
session create failed too busy	9	nas error
session failover protocol resync disconnect	6	admin reset
session hardware unavailable	8	port error
session no resources server port	9	nas error
session not ready	9	nas error
session rx cdn	10	nas request
session rx cdn avp bad hidden	10	nas request
session rx cdn avp bad value assigned session id	10	nas request
session rx cdn avp duplicate value assigned session id	10	nas request
session rx cdn avp malformed bad length	10	nas request
session rx cdn avp malformed truncated	10	nas request
session rx cdn avp missing mandatory assigned session id	10	nas request
session rx cdn avp missing mandatory result code	10	nas request
session rx cdn avp missing random vector	10	nas request
session rx cdn avp missing secret	10	nas request

Table 63: Default L2TP Mappings (*continued*)

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
session rx cdn avp unknown	10	nas request
session rx cdn no resources	10	nas request
session rx iccn avp bad hidden	10	nas request
session rx iccn avp bad value framing type	10	nas request
session rx iccn avp bad value proxy authen type	10	nas request
session rx iccn avp bad value unsupported proxy authen type	10	nas request
session rx iccn avp malformed bad length	10	nas request
session rx iccn avp malformed truncated	10	nas request
session rx iccn avp missing mandatory connect speed	10	nas request
session rx iccn avp missing mandatory framing type	10	nas request
session rx iccn avp missing mandatory proxy authen challenge	10	nas request
session rx iccn avp missing mandatory proxy authen id	10	nas request
session rx iccn avp missing mandatory proxy authen name	10	nas request
session rx iccn avp missing mandatory proxy authen response	10	nas request
session rx iccn avp missing random vector	10	nas request
session rx iccn avp missing secret	10	nas request
session rx iccn avp unknown	10	nas request
session rx iccn no resources	10	nas request
session rx iccn unexpected	10	nas request
session rx icrp avp bad hidden	10	nas request
session rx icrp avp bad value assigned session id	10	nas request
session rx icrp avp duplicate value assigned session id	10	nas request
session rx icrp avp malformed bad length	10	nas request

Table 63: Default L2TP Mappings (*continued*)

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
session rx icrp avp malformed truncated	10	nas request
session rx icrp avp missing mandatory assigned session id	10	nas request
session rx icrp avp missing random vector	10	nas request
session rx icrp avp missing secret	10	nas request
session rx icrp avp unknown	10	nas request
session rx icrp no resources	10	nas request
session rx icrp unexpected	10	nas request
session rx icrq admin close	6	admin reset
session rx icrq authenticate failed host	10	nas request
session rx icrq avp bad hidden	10	nas request
session rx icrq avp bad value assigned session id	10	nas request
session rx icrq avp bad value bearer type	10	nas request
session rx icrq avp bad value cisco nas port	10	nas request
session rx icrq avp duplicate value assigned session id	10	nas request
session rx icrq avp malformed bad length	10	nas request
session rx icrq avp malformed truncated	10	nas request
session rx icrq avp missing mandatory assigned session id	10	nas request
session rx icrq avp missing mandatory call serial number	10	nas request
session rx icrq avp missing random vector	10	nas request
session rx icrq avp missing secret	10	nas request
session rx icrq avp unknown	10	nas request
session rx icrq no resources	10	nas request
session rx icrq unexpected	10	nas request

Table 63: Default L2TP Mappings (*continued*)

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
session rx occn avp bad hidden	10	nas request
session rx occn avp bad value framing type	10	nas request
session rx occn avp malformed bad length	10	nas request
session rx occn avp malformed truncated	10	nas request
session rx occn avp missing mandatory connect speed	10	nas request
session rx occn avp missing mandatory framing type	10	nas request
session rx occn avp missing random vector	10	nas request
session rx occn avp missing secret	10	nas request
session rx occn avp unknown	10	nas request
session rx occn no resources	10	nas request
session rx occn unexpected	10	nas request
session rx ocrp avp bad hidden	10	nas request
session rx ocrp avp bad value assigned session id	10	nas request
session rx ocrp avp duplicate value assigned session id	10	nas request
session rx ocrp avp malformed bad length	10	nas request
session rx ocrp avp malformed truncated	10	nas request
session rx ocrp avp missing mandatory assigned session id	10	nas request
session rx ocrp avp missing random vector	10	nas request
session rx ocrp avp missing secret	10	nas request
session rx ocrp avp unknown	10	nas request
session rx ocrp no resources	10	nas request
session rx ocrp unexpected	10	nas request
session rx ocrq admin close	10	admin reset

Table 63: Default L2TP Mappings (*continued*)

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
session rx ocrq authenticate failed host	10	nas request
session rx ocrq avp bad hidden	10	nas request
session rx ocrq avp bad value assigned session id	10	nas request
session rx ocrq avp bad value bearer type	10	nas request
session rx ocrq avp bad value framing type	10	nas request
session rx ocrq avp duplicate value assigned session id	10	nas request
session rx ocrq avp malformed bad length	10	nas request
session rx ocrq avp malformed truncated	10	nas request
session rx ocrq avp missing mandatory assigned session id	10	nas request
session rx ocrq avp missing mandatory bearer type	10	nas request
session rx ocrq avp missing mandatory call serial number	10	nas request
session rx ocrq avp missing mandatory called number	10	nas request
session rx ocrq avp missing mandatory framing type	10	nas request
session rx ocrq avp missing mandatory maximum bps	10	nas request
session rx ocrq avp missing mandatory minimum bps	10	nas request
session rx ocrq avp missing random vector	10	nas request
session rx ocrq avp missing secret	10	nas request
session rx ocrq avp unknown	10	nas request
session rx ocrq no resources	10	nas request
session rx ocrq unexpected	10	nas request
session rx ocrq unsupported	9	nas error
session rx sli avp bad hidden	10	nas request
session rx sli avp bad value accm	10	nas request

Table 63: Default L2TP Mappings (*continued*)

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
session rx sli avp malformed bad length	10	nas request
session rx sli avp malformed truncated	10	nas request
session rx sli avp missing mandatory accm	10	nas request
session rx sli avp missing random vector	10	nas request
session rx sli avp missing secret	10	nas request
session rx sli avp unknown	10	nas request
session rx sli no resources	10	nas request
session rx unexpected packet lac incoming	10	nas request
session rx unexpected packet lac outgoing	10	nas request
session rx unexpected packet lns incoming	10	nas request
session rx unexpected packet lns outgoing	10	nas request
session rx unknown session id	10	nas request
session rx wen avp bad hidden	10	nas request
session rx wen avp malformed bad length	10	nas request
session rx wen avp malformed truncated	10	nas request
session rx wen avp missing mandatory call errors	10	nas request
session rx wen avp missing random vector	10	nas request
session rx wen avp missing secret	10	nas request
session rx wen avp unknown	10	nas request
session rx wen no resources	10	nas request
session timeout connection	10	nas request
session timeout inactivity	4	idle timeout
session timeout session	5	session timeout

Table 63: Default L2TP Mappings (*continued*)

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
session timeout upper create	9	nas error
session transmit speed unavailable	9	nas error
session tunnel down	15	service unavailable
session tunnel failed	15	service unavailable
session tunnel switch profile deleted	6	admin reset
session tunneled interface down	8	port error
session unknown cause	9	nas error
session upper create failed	9	nas error
session upper removed	15	service unavailable
session warmstart not operational	15	service unavailable
session warmstart recovery error	15	service unavailable
session warmstart upper not restacked	10	nas request
tunnel admin close	6	admin reset
tunnel admin drain	6	admin reset
tunnel control channel failed	15	service unavailable
tunnel created no sessions	1	user request
tunnel destination address changed	6	admin reset
tunnel destination down	10	nas request
tunnel failover protocol no resources for recovery tunnel	15	service unavailable
tunnel failover protocol no resources for session resync	15	service unavailable
tunnel failover protocol not supported	15	service unavailable
tunnel failover protocol not supported by peer	15	service unavailable
tunnel failover protocol recovery control channel failed	15	service unavailable

Table 63: Default L2TP Mappings (*continued*)

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
tunnel failover protocol recovery tunnel failed	15	service unavailable
tunnel failover protocol recovery tunnel finished	1	user request
tunnel failover protocol recovery tunnel primary down	1	user request
tunnel failover protocol session resync failed	15	service unavailable
tunnel host profile changed	6	admin reset
tunnel host profile deleted	6	admin reset
tunnel rx scccn authenticate failed challenge	17	user error
tunnel rx scccn avp bad hidden	15	service unavailable
tunnel rx scccn avp bad value challenge response	15	service unavailable
tunnel rx scccn avp malformed bad length	15	service unavailable
tunnel rx scccn avp malformed truncated	15	service unavailable
tunnel rx scccn avp missing challenge response	17	user error
tunnel rx scccn avp missing random vector	15	service unavailable
tunnel rx scccn avp missing secret	15	service unavailable
tunnel rx scccn avp unexpected challenge response	15	service unavailable
tunnel rx scccn avp unknown	15	service unavailable
tunnel rx scccn no resources	15	service unavailable
tunnel rx scccn session id not null	15	service unavailable
tunnel rx scccn unexpected	15	service unavailable
tunnel rx sccrp authenticate failed challenge	17	user error
tunnel rx sccrp authenticate failed host	17	user error
tunnel rx sccrp avp bad hidden	15	service unavailable
tunnel rx sccrp avp bad value assigned tunnel id	15	service unavailable

Table 63: Default L2TP Mappings (*continued*)

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
tunnel rx sccrp avp bad value bearer capabilities	15	service unavailable
tunnel rx sccrp avp bad value challenge	15	service unavailable
tunnel rx sccrp avp bad value challenge response	15	service unavailable
tunnel rx sccrp avp bad value failover capability	15	service unavailable
tunnel rx sccrp avp bad value framing capabilities	15	service unavailable
tunnel rx sccrp avp bad value protocol version	15	service unavailable
tunnel rx sccrp avp bad value receive window size	15	service unavailable
tunnel rx sccrp avp duplicate value assigned tunnel id	15	service unavailable
tunnel rx sccrp avp malformed bad length	15	service unavailable
tunnel rx sccrp avp malformed truncated	15	service unavailable
tunnel rx sccrp avp missing challenge response	17	user error
tunnel rx sccrp avp missing mandatory assigned tunnel id	15	service unavailable
tunnel rx sccrp avp missing mandatory framing capabilities	15	service unavailable
tunnel rx sccrp avp missing mandatory host name	15	service unavailable
tunnel rx sccrp avp missing mandatory protocol version	15	service unavailable
tunnel rx sccrp avp missing random vector	15	service unavailable
tunnel rx sccrp avp missing secret	15	service unavailable
tunnel rx sccrp avp unexpected challenge response	15	service unavailable
tunnel rx sccrp avp unexpected challenge without secret	15	service unavailable
tunnel rx sccrp avp unknown	15	service unavailable
tunnel rx sccrp no resources	15	service unavailable
tunnel rx sccrp session id not null	15	service unavailable
tunnel rx sccrp unexpected	15	service unavailable

Table 63: Default L2TP Mappings (*continued*)

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
tunnel rx sccrq admin close	6	admin reset
tunnel rx sccrq authenticate failed host	17	user error
tunnel rx sccrq avp bad hidden	15	service unavailable
tunnel rx sccrq avp bad value assigned tunnel id	15	service unavailable
tunnel rx sccrq avp bad value bearer capabilities	15	service unavailable
tunnel rx sccrq avp bad value challenge	15	service unavailable
tunnel rx sccrq avp bad value failover capability	15	service unavailable
tunnel rx sccrq avp bad value framing capabilities	15	service unavailable
tunnel rx sccrq avp bad value protocol version	15	service unavailable
tunnel rx sccrq avp bad value receive window size	15	service unavailable
tunnel rx sccrq avp duplicate value assigned tunnel id	15	service unavailable
tunnel rx sccrq avp malformed bad length	15	service unavailable
tunnel rx sccrq avp malformed truncated	15	service unavailable
tunnel rx sccrq avp missing mandatory assigned tunnel id	15	service unavailable
tunnel rx sccrq avp missing mandatory framing capabilities	15	service unavailable
tunnel rx sccrq avp missing mandatory host name	15	service unavailable
tunnel rx sccrq avp missing mandatory protocol version	15	service unavailable
tunnel rx sccrq avp missing random vector	15	service unavailable
tunnel rx sccrq avp missing secret	15	service unavailable
tunnel rx sccrq avp unexpected challenge without secret	15	service unavailable
tunnel rx sccrq avp unknown	15	service unavailable
tunnel rx sccrq bad address	15	service unavailable
tunnel rx sccrq no resources	15	service unavailable

Table 63: Default L2TP Mappings (*continued*)

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
tunnel rx sccrq no resources max tunnels	15	service unavailable
tunnel rx sccrq session id not null	15	service unavailable
tunnel rx sccrq unexpected	15	service unavailable
tunnel rx stopccn	1	user request
tunnel rx stopccn avp bad hidden	15	service unavailable
tunnel rx stopccn avp bad value assigned tunnel id	15	service unavailable
tunnel rx stopccn avp duplicate value assigned tunnel id	15	service unavailable
tunnel rx stopccn avp malformed bad length	15	service unavailable
tunnel rx stopccn avp malformed truncated	15	service unavailable
tunnel rx stopccn avp missing mandatory assigned tunnel id	15	service unavailable
tunnel rx stopccn avp missing mandatory result code	15	service unavailable
tunnel rx stopccn avp missing random vector	15	service unavailable
tunnel rx stopccn avp missing secret	15	service unavailable
tunnel rx stopccn avp unknown	15	service unavailable
tunnel rx stopccn no resources	15	service unavailable
tunnel rx stopccn session id not null	15	service unavailable
tunnel rx frs avp malformed truncated	15	service unavailable
tunnel rx frs avp missing mandatory failover session state	15	service unavailable
tunnel rx frs avp missing random vector	15	service unavailable
tunnel rx frs avp missing secret	15	service unavailable
tunnel rx frs avp unknown	15	service unavailable
tunnel rx frs no resources	15	service unavailable
tunnel rx frs session id not null	15	service unavailable

Table 63: Default L2TP Mappings (*continued*)

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
tunnel rx fsq avp bad hidden	15	service unavailable
tunnel rx fsq avp malformed bad length	15	service unavailable
tunnel rx fsq avp malformed truncated	15	service unavailable
tunnel rx fsq avp missing mandatory failover session state	15	service unavailable
tunnel rx fsq avp missing random vector	15	service unavailable
tunnel rx fsq avp missing secret	15	service unavailable
tunnel rx fsq avp unknown	15	service unavailable
tunnel rx fsq no resources	15	service unavailable
tunnel rx fsq session id not null	15	service unavailable
tunnel rx fsr avp bad hidden	15	service unavailable
tunnel rx fsr avp malformed bad length	15	service unavailable
tunnel rx unexpected packet	15	service unavailable
tunnel rx unexpected packet for session	15	service unavailable
tunnel rx unknown packet message type indecipherable	15	service unavailable
tunnel rx unknown packet message type unrecognized	15	service unavailable
tunnel rx recovery sccn authenticate failed challenge	17	user error
tunnel rx recovery sccn avp bad hidden	15	service unavailable
tunnel rx recovery sccn avp bad value challenge response	15	service unavailable
tunnel rx recovery sccn avp malformed bad length	15	service unavailable
tunnel rx recovery sccn avp malformed truncated	15	service unavailable
tunnel rx recovery sccn avp missing challenge response	17	user error
tunnel rx recovery sccn avp missing random vector	15	service unavailable
tunnel rx recovery sccn avp missing secret	15	service unavailable

Table 63: Default L2TP Mappings (*continued*)

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
tunnel rx recovery scccn avp unexpected challenge response	15	service unavailable
tunnel rx recovery scccn avp unknown	15	service unavailable
tunnel rx recovery scccn no resources	15	service unavailable
tunnel rx recovery scccn session id not null	15	service unavailable
tunnel rx recovery sccrp authenticate failed challenge	17	user error
tunnel rx recovery sccrp avp bad hidden	15	service unavailable
tunnel rx recovery sccrp avp bad value assigned tunnel id	15	service unavailable
tunnel rx recovery sccrp avp bad value bearer capabilities	15	service unavailable
tunnel rx recovery sccrp avp bad value challenge	15	service unavailable
tunnel rx recovery sccrp avp bad value challenge response	15	service unavailable
tunnel rx recovery sccrp avp bad value framing capabilities	15	service unavailable
tunnel rx recovery sccrp avp bad value protocol version	15	service unavailable
tunnel rx recovery sccrp avp bad value receive window size	15	service unavailable
tunnel rx recovery sccrp avp bad value suggested control sequence	15	service unavailable
tunnel rx recovery sccrp avp duplicate value assigned tunnel id	15	service unavailable
tunnel rx recovery sccrp avp malformed bad length	15	service unavailable
tunnel rx recovery sccrp avp malformed truncated	15	service unavailable
tunnel rx recovery sccrp avp mismatched host name	15	service unavailable
tunnel rx recovery sccrp avp mismatched vendor name	15	service unavailable
tunnel rx recovery sccrp avp missing challenge response	17	user error
tunnel rx recovery sccrp avp missing mandatory assigned tunnel id	15	service unavailable

Table 63: Default L2TP Mappings (*continued*)

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
tunnel rx recovery sccrp avp missing mandatory framing capabilities	15	service unavailable
tunnel rx recovery sccrp avp missing mandatory host name	15	service unavailable
tunnel rx recovery sccrp avp missing mandatory protocol version	15	service unavailable
tunnel rx recovery sccrp avp missing random vector	15	service unavailable
tunnel rx recovery sccrp avp missing secret	15	service unavailable
tunnel rx recovery sccrp avp unexpected challenge response	15	service unavailable
tunnel rx recovery sccrp avp unexpected challenge without secret	15	service unavailable
tunnel rx recovery sccrp avp unknown	15	service unavailable
tunnel rx recovery sccrp no resources	15	service unavailable
tunnel rx recovery sccrp session id not null	15	service unavailable
tunnel rx recovery sccrq admin close	6	admin reset
tunnel rx recovery sccrq avp bad hidden	15	service unavailable
tunnel rx recovery sccrq avp bad value assigned tunnel id	15	service unavailable
tunnel rx recovery sccrq avp bad value bearer capabilities	15	service unavailable
tunnel rx recovery sccrq avp bad value challenge	15	service unavailable
tunnel rx recovery sccrq avp bad value framing capabilities	15	service unavailable
tunnel rx recovery sccrq avp bad value protocol version	15	service unavailable
tunnel rx recovery sccrq avp bad value receive window size	15	service unavailable
tunnel rx recovery sccrq avp bad value tunnel recovery	15	service unavailable
tunnel rx recovery sccrq avp duplicate value assigned tunnel id	15	service unavailable
tunnel rx recovery sccrq avp duplicate value tie breaker	15	service unavailable

Table 63: Default L2TP Mappings (*continued*)

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
tunnel rx recovery sccrq avp malformed bad length	15	service unavailable
tunnel rx recovery sccrq avp malformed truncated	15	service unavailable
tunnel rx recovery sccrq avp mismatched host name	15	service unavailable
tunnel rx recovery sccrq avp mismatched vendor name	15	service unavailable
tunnel rx recovery sccrq avp missing mandatory assigned tunnel id	15	service unavailable
tunnel rx recovery sccrq avp missing mandatory framing capabilities	15	service unavailable
tunnel rx recovery sccrq avp missing mandatory host name	15	service unavailable
tunnel rx recovery sccrq avp missing mandatory protocol version	15	service unavailable
tunnel rx recovery sccrq avp missing mandatory tunnel recovery	15	service unavailable
tunnel rx recovery sccrq avp missing random vector	15	service unavailable
tunnel rx recovery sccrq avp missing secret	15	service unavailable
tunnel rx recovery sccrq avp missing tie breaker	15	service unavailable
tunnel rx recovery sccrq avp unexpected challenge without secret	15	service unavailable
tunnel rx recovery sccrq avp unknown	15	service unavailable
tunnel rx recovery sccrq no resources	15	service unavailable
tunnel rx recovery sccrq session id not null	15	service unavailable
tunnel rx recovery sccrq tunnel id not null	15	service unavailable
tunnel rx recovery stopccn avp bad hidden	15	service unavailable
tunnel rx recovery stopccn avp bad value assigned tunnel id	15	service unavailable
tunnel rx recovery stopccn avp duplicate value assigned tunnel id	15	service unavailable

Table 63: Default L2TP Mappings (*continued*)

L2TP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
tunnel rx recovery stopccn avp malformed bad length	15	service unavailable
tunnel rx recovery stopccn avp malformed truncated	15	service unavailable
tunnel rx recovery stopccn avp missing mandatory assigned tunnel id	15	service unavailable
tunnel rx recovery stopccn avp missing mandatory result code	15	service unavailable
tunnel rx recovery stopccn avp missing random vector	15	service unavailable
tunnel rx recovery stopccn avp missing secret	15	service unavailable
tunnel rx recovery stopccn avp unknown	15	service unavailable
tunnel rx recovery stopccn no resources	15	service unavailable
tunnel rx recovery stopccn session id not null	15	service unavailable
tunnel rx recovery unexpected packet	15	service unavailable
tunnel rx recovery unknown packet message type indecipherable	15	service unavailable
tunnel rx recovery unknown packet message type unrecognized	15	service unavailable
tunnel rx session packet null sid invalid	15	service unavailable
tunnel rx session packet null sid without assigned session id	15	service unavailable
tunnel timeout connection	15	service unavailable
tunnel timeout connection recovery tunnel	15	service unavailable
tunnel timeout idle	1	user request
tunnel unknown cause	9	nas error
tunnel warmstart not operational	15	service unavailable
tunnel warmstart recovery error	15	service unavailable

- Related Documentation**
- [Mapping Application Terminate Reasons and RADIUS Terminate Codes Overview on page 36](#)
 - [Configuring Custom Mappings for PPP Terminate Reasons on page 95](#)
 - [Monitoring Application Terminate Reason Mappings on page 167](#)

PPP Terminate Reasons

Table 64 on page 283 lists the default PPP terminate mappings. The table indicates the supported PPP terminate reasons and the RADIUS Acct-Terminate-Cause attributes they are mapped to by default.

Table 64: Default PPP Mappings

PPP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
authenticate authenticator timeout	17	user error
authenticate challenge timeout	10	nas request
authenticate chap no resources	10	nas request
authenticate chap peer authenticator timeout	17	user error
authenticate deny by peer	17	user error
authenticate inactivity timeout	4	idle timeout
authenticate max requests	10	nas request
authenticate no authenticator	10	nas request
authenticate pap peer authenticator timeout	17	user error
authenticate pap request timeout	10	nas request
authenticate session timeout	5	session timeout
authenticate too many requests	10	nas request
authenticate tunnel fail immediate	10	nas request
authenticate tunnel unsupported tunnel type	10	nas request
bundle fail create	10	nas request
bundle fail engine add	10	nas request

Table 64: Default PPP Mappings (*continued*)

PPP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
bundle fail fragment size mismatch	10	nas request
bundle fail fragmentation location	10	nas request
bundle fail fragmentation mismatch	10	nas request
bundle fail join	10	nas request
bundle fail link selection mismatch	10	nas request
bundle fail local mped not set yet	10	nas request
bundle fail local mrru mismatch	10	nas request
bundle fail local mru mismatch	10	nas request
bundle fail peer mrru mismatch	10	nas request
bundle fail reassembly location	10	nas request
bundle fail reassembly mismatch	10	nas request
bundle fail record network	10	nas request
bundle fail server location mismatch	10	nas request
bundle fail static link	10	nas request
failover during authentication	6	admin reset
interface admin disable	6	admin reset
interface down	2	lost carrier
interface no hardware	8	port error
ip admin disable	10	nas request
ip inhibited by authentication	10	nas request
ip link down	10	nas request
ip max configure exceeded	10	nas request
ip no local ip address	10	nas request

Table 64: Default PPP Mappings (continued)

PPP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
ip no local ip address mask	10	nas request
ip no local primary dns address	10	nas request
ip no local primary nbns address	10	nas request
ip no local secondary dns address	10	nas request
ip no local secondary nbns address	10	nas request
ip no peer ip address	10	nas request
ip no peer ip address mask	10	nas request
ip no peer primary dns address	10	nas request
ip no peer primary nbns address	10	nas request
ip no peer secondary dns address	10	nas request
ip no peer secondary nbns address	10	nas request
ip no service	10	nas request
ip peer renegotiate rx conf ack	10	nas request
ip peer renegotiate rx conf nak	10	nas request
ip peer renegotiate rx conf rej	10	nas request
ip peer renegotiate rx conf req	10	nas request
ip peer terminate term ack	10	nas request
ip peer terminate code rej	10	nas request
ip peer terminate term req	10	nas request
ip service disable	10	nas request
ip stale stacking	10	nas request
ipv6 admin disable	10	nas request
ipv6 inhibited by authentication	10	nas request

Table 64: Default PPP Mappings (*continued*)

PPP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
ipv6 link down	10	nas request
ipv6 local and peer interface ids identical	10	nas request
ipv6 max configure exceeded	10	nas request
ipv6 no local ipv6 interface id	10	nas request
ipv6 no peer ipv6 interface id	10	nas request
ipv6 no service	10	nas request
ipv6 peer renegotiate rx conf ack	10	nas request
ipv6 peer renegotiate rx conf nak	10	nas request
ipv6 peer renegotiate rx conf rej	10	nas request
ipv6 peer renegotiate rx conf req	10	nas request
ipv6 peer terminate code rej	10	nas request
ipv6 peer terminate term ack	10	nas request
ipv6 peer terminate term req	10	nas request
ipv6 service disable	10	nas request
ipv6 stale stacking	10	nas request
lcp authenticate terminate hold	10	nas request
lcp configured mrru too small	10	nas request
lcp configured mru invalid	10	nas request
lcp configured mru too small	10	nas request
lcp dynamic interface hold	10	nas request
lcp keepalive failure	10	nas request
lcp loopback rx conf req	10	nas request
lcp loopback rx echo reply	10	nas request

Table 64: Default PPP Mappings (*continued*)

PPP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
lcp loopback rx echo req	10	nas request
lcp max configure exceeded	10	nas request
lcp mru changed	10	nas request
lcp negotiation timeout	10	nas request
lcp no localacm	10	nas request
lcp no localacfc	10	nas request
lcp no local authentication	10	nas request
lcp no local endpoint discriminator	10	nas request
lcp no local magic number	10	nas request
lcp no local mrru	10	nas request
lcp no local mru	10	nas request
lcp no localpfc	10	nas request
lcp no peer accm	10	nas request
lcp no peer authentication	10	nas request
lcp no peer endpoint discriminator	10	nas request
lcp no peer magicnumber	10	nas request
lcp no peer mrru	10	nas request
lcp no peer mru	10	nas request
lcp no peer pfc	10	nas request
lcp peer terminate code rej	1	user request
lcp peer terminate term ack	1	user request
lcp peer terminate term req	1	user request
lcp peer terminate protocol reject	1	user request

Table 64: Default PPP Mappings (*continued*)

PPP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
lcp peer renegotiate rx conf ack	1	user request
lcp peer renegotiate rx conf nak	1	user request
lcp peer renegotiate rx conf rej	1	user request
lcp peer renegotiate rx conf req	1	user request
lcp tunnel disconnected	10	nas request
lcp tunnel failed	10	nas request
link interface no hardware	8	port error
lower interface attach failed	2	lost carrier
lower interface teardown	2	lost carrier
mpls admin disable	10	nas request
mpls link down	10	nas request
mpls max configure exceeded	10	nas request
mpls no service	10	nas request
mpls peer renegotiate rx conf ack	10	nas request
mpls peer renegotiate rx conf nak	10	nas request
mpls peer renegotiate rx conf rej	10	nas request
mpls peer renegotiate rx conf req	10	nas request
mpls peer terminate code rej	10	nas request
mpls peer terminate term ack	10	nas request
mpls peer terminate term req	10	nas request
mpls service disable	10	nas request
mpls stale stacking	10	nas request
network interface admin disable	6	admin reset

Table 64: Default PPP Mappings (continued)

PPP Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
no bundle	10	nas request
no interface	8	port error
no link interface	8	port error
no ncps available	10	nas request
no network interface	10	nas request
no upper interface	9	nas error
osi admin disable	10	nas request
osi link down	10	nas request
osi max configure exceeded	10	nas request
osi no local align npdu	10	nas request
osi no peer align npdu	10	nas request
osi no service	10	nas request
osi peer renegotiate rx conf ack	10	nas request
osi peer renegotiate rx conf nak	10	nas request
osi peer renegotiate rx conf rej	10	nas request
osi peer renegotiate rx conf req	10	nas request
osi peer terminate code rej	10	nas request
osi peer terminate term ack	10	nas request
osi peer terminate term req	10	nas request
osi service disable	10	nas request
osi stale stacking	10	nas request

**Related
Documentation**

- [Mapping Application Terminate Reasons and RADIUS Terminate Codes Overview on page 36](#)

- [Configuring Custom Mappings for PPP Terminate Reasons on page 95](#)
- [L2TP Terminate Reasons on page 266](#)
- [Monitoring Application Terminate Reason Mappings on page 167](#)

RADIUS Client Terminate Reasons

[Table 65 on page 290](#) lists the default RADIUS client terminate mappings. The table indicates the supported RADIUS client terminate reasons and the RADIUS Acct-Terminate-Cause attributes they are mapped to by default.

Table 65: Default RADIUS Client Mappings

RADIUS Client Terminate Reason	RADIUS Acct-Terminate-Cause	
	Code	Description
no-acct-server	10	nas request
system-reboot	10	nas request
virtual-router-deletion	10	nas request

Related Documentation

- [Mapping Application Terminate Reasons and RADIUS Terminate Codes Overview on page 36](#)
- [Monitoring Application Terminate Reason Mappings on page 167](#)

CHAPTER 9

Monitoring RADIUS

This chapter describes how to monitor the RADIUS attributes, RADIUS dynamic-request server, and RADIUS relay.

RADIUS topics are described in the following sections:

- [Monitoring Override Settings of RADIUS IETF Attributes on page 291](#)
- [Monitoring the NAS-Port-Format RADIUS Attribute on page 292](#)
- [Monitoring the Calling-Station-Id RADIUS Attribute on page 293](#)
- [Monitoring the NAS-Identifier RADIUS Attribute on page 293](#)
- [Monitoring the Format of the Remote-Circuit-ID for RADIUS on page 294](#)
- [Monitoring the Delimiter Character in the Remote-Circuit-ID for RADIUS on page 294](#)
- [Monitoring the Acct-Session-Id RADIUS Attribute on page 294](#)
- [Monitoring the DSL-Port-Type RADIUS Attribute on page 294](#)
- [Monitoring the Connect-Info RADIUS Attribute on page 295](#)
- [Monitoring the NAS-Port-ID RADIUS Attribute on page 295](#)
- [Monitoring Included RADIUS Attributes on page 295](#)
- [Monitoring Ignored RADIUS Attributes on page 298](#)
- [Setting the Baseline for RADIUS Dynamic-Request Server Statistics on page 298](#)
- [Monitoring RADIUS Dynamic-Request Server Statistics on page 299](#)
- [Monitoring the Configuration of the RADIUS Dynamic-Request Server on page 300](#)
- [Setting a Baseline for RADIUS Relay Statistics on page 301](#)
- [Monitoring RADIUS Relay Server Statistics on page 301](#)
- [Monitoring the Configuration of the RADIUS Relay Server on page 303](#)
- [Monitoring the Status of RADIUS Relay UDP Checksums on page 303](#)
- [Monitoring the Status of ICR Partition Accounting on page 304](#)

Monitoring Override Settings of RADIUS IETF Attributes

Purpose Display the current override setting for RADIUS IETF attributes. You can monitor the NAS-IP-Address [4], NAS-Port-Id [87], Calling-Station-Id [31], and NAS-Identifier [32] attributes.

Action To display the current setting for all configured RADIUS attributes:

```
host1#show radius override
nas-ip-addr:      nas-ip-addr
nas-port-id:      nas-port-id
calling-station-id: calling-station-id
nas-info:         from current virtual router

host1#show radius override
nas-ip-addr: nas-ip-addr
nas-info:    from authentication virtual router
```

Meaning [Table 66 on page 292](#) lists the **show radius override** command output fields.

Table 66: show radius override Output Fields

Field Name	Field Description
nas-ip-addr	Displays the current setting for the NAS-IP-Address [4] attribute. These settings can be changed with the radius override nas-ip-addr tunnel-client-endpoint and radius override nas-info commands.
nas-port-id	Displays the current setting for the NAS-Port-Id [87] attribute. Use the radius override nas-port-id remote-circuit-id command to override the standard NAS-Port-Id attribute with the PPPoE remote circuit ID transmitted from the DSLAM.
calling-station-id	Displays the current setting for the Calling-Station-Id [31] attribute. Use the radius override calling-station-id remote-circuit-id command to override the standard Calling-Station-Id attribute with the PPPoE remote circuit ID transmitted from the DSLAM.
nas-info	Displays the current setting for the NAS-Identifier [32] attribute. This setting can be changed with the radius override nas-info command, which is used for AAA broadcast accounting.

Related Documentation

- *show radius override*

Monitoring the NAS-Port-Format RADIUS Attribute

Purpose Display information for the NAS-Port attribute.

Action To display the setting for the NAS-Port attribute:

```
host1#show radius nas-port-format
0ssssppp
```

To display information about the NAS-Port attribute on an ATM interface on an E320 Broadband Services Router:

```
host1#show radius nas-port-format extended atm
extended atm field-width slot 5 adapter 0 port 4 vpi 4 vci 12
```

To display the status of NAS-Port attribute settings for PPPoE interfaces:

```
host1#show radius pppoe nas-port-format
unique
```

To display the status of the S-VLAN ID setting for the NAS-Port attribute for VLAN interfaces:

```
host1#show radius vlan nas-port-format
vlan stacked
```

- Related Documentation**
- *show radius nas-port-format*
 - *show radius nas-port-format extended*
 - *show radius pppoe nas-port-format*
 - *show radius vlan nas-port-format*

Monitoring the Calling-Station-Id RADIUS Attribute

Purpose Display the format and delimiter used for the Calling-Station-Id [31] attribute.

Action To display the format configured for the Calling-Station-Id [31] attribute:

```
host1#show radius calling-station-format
fixed-format-adapter-new-field (includes SVLAN ID)
```

To display the delimiter used in the Calling-Station-Id for authenticated ATM PPP users:

```
host1#show radius calling-station-delimiter
&
```

- Related Documentation**
- *show radius calling-station-format*
 - *show radius calling-station-delimiter*

Monitoring the NAS-Identifier RADIUS Attribute

Purpose Display information about the NAS-Identifier value.

Action To display information about the NAS-Identifier value:

```
host1#show radius nas-identifier
fox
```

- Related Documentation**
- *show radius nas-identifier*

Monitoring the Format of the Remote-Circuit-ID for RADIUS

- Purpose** Display the format configured for the PPPoE remote circuit ID value captured from a DSLAM.
- The default format is agent-circuit-ID. If the PPPoE remote circuit ID value is configured to include any or all of the agent-circuit-id, agent-remote-id, and nas-identifier components, the display lists the components included and the order in which they appear.
- If the PPPoE remote circuit ID value is configured to use the format for the **dsl-forum-1** keyword of **radius remote-circuit-id-format**, the display indicates that this format is in effect.
- Action** To display the format configured for the PPPoE remote circuit ID value captured from a DSLAM:
- ```
host1#show radius remote-circuit-id-format
nas-identifier agent-circuit-id agent-remote-id
```

**Related Documentation**

- *show radius remote-circuit-id-format*

## Monitoring the Delimiter Character in the Remote-Circuit-ID for RADIUS

---

- Purpose** Display the delimiter character configured to set off components in the PPPoE remote circuit ID value captured from a DSLAM. The default delimiter character is #.
- Action** To display the delimiter character:
- ```
host1#show radius remote-circuit-id-delimiter
!
```

Related Documentation

- *show radius remote-circuit-id-delimiter*

Monitoring the Acct-Session-Id RADIUS Attribute

- Purpose** Display the format used for the Acct-Session-Id attribute.
- Action** To display the format used for the Acct-Session-Id attribute:
- ```
host1#show radius acct-session-id-format
decimal
```

**Related Documentation**

- *show radius acct-session-id-format*

## Monitoring the DSL-Port-Type RADIUS Attribute

---

- Purpose** Display the DSL port type for NAS-Port-Type attribute for ATM and Ethernet users.

**Action** To display the DSL port type for NAS-Port-Type attribute for ATM users:

```
host1#show radius dsl-port-type
xds1
```

To display the NAS-Port-Type attribute for Ethernet interfaces:

```
host1#show radius ethernet-port-type
virtual
```

**Related Documentation**

- *show radius dsl-port-type*
- *show radius ethernet-port-type*

## Monitoring the Connect-Info RADIUS Attribute

**Purpose** Display the format for the Connect-Info attribute.

**Action** To display the format for the Connect-Info attribute:

```
host1(config)#show radius connect-info-format
12tp-connect-speed-rx-when-equal
```

**Related Documentation**

- *show radius connect-info-format*

## Monitoring the NAS-Port-ID RADIUS Attribute

**Purpose** Display whether the router includes or excludes the subinterface number or adapter in the interface description that the router passes to RADIUS for inclusion in the NAS-Port-Id attribute.

**Action** To display information about the interface description for the NAS-Port-ID:

```
host1#show aaa intf-desc-format
exclude sub-interface
include adapter
```

**Related Documentation**

- *show aaa intf-desc-format*

## Monitoring Included RADIUS Attributes

**Purpose** Display the RADIUS attributes that are included in and excluded from Acct-On, Acct-Off, Access-Request, Acct-Start, and Acct-Stop messages.

**Action** To display the list of included RADIUS attributes:

```
host1# show radius attributes-included
```

| Attribute Name         | Account On | Account Off | Access Request | Account Start | Account Stop |
|------------------------|------------|-------------|----------------|---------------|--------------|
| access-loop-parameters | n/c        | n/c         | disabled       | n/c           | n/c          |
| acct-authentic         | enabled    | enabled     | n/c            | n/c           | n/c          |
| acct-delay-time        | enabled    | enabled     | n/c            | n/c           | n/c          |

|                                |         |         |          |          |          |
|--------------------------------|---------|---------|----------|----------|----------|
| acct-link-count                | n/c     | n/c     | n/c      | enabled  | enabled  |
| acct-multi-session-id          | n/c     | n/c     | disabled | enabled  | enabled  |
| acct-session-id                | enabled | enabled | enabled  | n/c      | n/c      |
| acct-terminate-cause           | n/c     | enabled | n/c      | n/c      | n/c      |
| acct-tunnel-connection         | n/c     | n/c     | enabled  | enabled  | enabled  |
| ascend-num-in-multilink        | n/c     | n/c     | disabled | disabled | disabled |
| called-station-id              | n/c     | n/c     | enabled  | enabled  | enabled  |
| downstream-calculated-qos-rate | n/c     | n/c     | disabled | disabled | disabled |
| upstream-calculated-qos-rate   | n/c     | n/c     | disabled | disabled | disabled |
| calling-station-id             | n/c     | n/c     | enabled  | enabled  | enabled  |
| class                          | n/c     | n/c     | n/c      | enabled  | enabled  |
| connect-info                   | n/c     | n/c     | enabled  | enabled  | enabled  |
| delegated-ipv6-prefix          | n/c     | n/c     | n/c      | disabled | disabled |
| dhcp-options                   | n/c     | n/c     | disabled | disabled | disabled |
| dhcp-option-82                 | n/c     | n/c     | disabled | disabled | disabled |
| dhcp-option82-circuitid        | n/c     | n/c     | enabled  | disabled | disabled |
| dhcp-option82-remoteid         | n/c     | n/c     | enabled  | disabled | disabled |
| dhcp-mac-address               | n/c     | n/c     | disabled | disabled | disabled |
| dhcp-gi-address                | n/c     | n/c     | disabled | disabled | disabled |
| dsl-forum-attributes(vsa)      | n/c     | n/c     | disabled | disabled | disabled |
| ds-lite-tunnel-name            | n/c     | n/c     | n/c      | disabled | disabled |
| egress-policy-name(vsa)        | n/c     | n/c     | n/c      | enabled  | enabled  |
| ipv6-egress-policy-name(vsa)   | n/c     | n/c     | n/c      | enabled  | enabled  |
| event-timestamp                | enabled | enabled | n/c      | enabled  | enabled  |
| framed-compression             | n/c     | n/c     | n/c      | enabled  | enabled  |
| framed-interface-id            | n/c     | n/c     | n/c      | disabled | disabled |
| framed-ip-address              | n/c     | n/c     | n/c      | enabled  | enabled  |
| framed-ip-netmask              | n/c     | n/c     | n/c      | enabled  | enabled  |
| framed-ipv6-pool               | n/c     | n/c     | n/c      | disabled | disabled |
| framed-ipv6-prefix             | n/c     | n/c     | n/c      | disabled | disabled |
| framed-ipv6-route              | n/c     | n/c     | n/c      | disabled | disabled |
| framed-route                   | n/c     | n/c     | n/c      | disabled | disabled |
| icr-partition-id(vsa)          | n/c     | n/c     | disabled | disabled | disabled |
| ingress-policy-name(vsa)       | n/c     | n/c     | n/c      | enabled  | enabled  |
| ipv6-ingress-policy-name(vsa)  | n/c     | n/c     | n/c      | enabled  | enabled  |
| input-gigapkts(vsa)            | n/c     | n/c     | n/c      | n/c      | enabled  |
| input-gigawords                | n/c     | n/c     | n/c      | n/c      | enabled  |
| interface-description(vsa)     | n/c     | n/c     | enabled  | enabled  | enabled  |
| ipv6-input-octets(vsa)         | n/c     | n/c     | n/c      | n/c      | enabled  |
| ipv6-output-octets(vsa)        | n/c     | n/c     | n/c      | n/c      | enabled  |
| ipv6-input-packets(vsa)        | n/c     | n/c     | n/c      | n/c      | enabled  |
| ipv6-output-packets(vsa)       | n/c     | n/c     | n/c      | n/c      | enabled  |
| ipv6-input-gigawords(vsa)      | n/c     | n/c     | n/c      | n/c      | enabled  |
| ipv6-output-gigawords(vsa)     | n/c     | n/c     | n/c      | n/c      | enabled  |
| ipv6-local-interface(vsa)      | n/c     | n/c     | n/c      | disabled | disabled |
| ipv6-nd-ra-prefix(vsa)         | n/c     | n/c     | n/c      | disabled | disabled |
| ipv6-primary-dns(vsa)          | n/c     | n/c     | n/c      | disabled | disabled |
| ipv6-secondary-dns(vsa)        | n/c     | n/c     | n/c      | disabled | disabled |
| ipv6-virtual-router(vsa)       | n/c     | n/c     | n/c      | disabled | disabled |
| l2c-downstream-data(vsa)       | n/c     | n/c     | disabled | disabled | disabled |
| l2c-upstream-data(vsa)         | n/c     | n/c     | disabled | disabled | disabled |
| l2cd-acc-loop-cir-id(vsa)      | n/c     | n/c     | disabled | disabled | disabled |
| l2cd-acc-aggr-cir-id-bin(vsa)  | n/c     | n/c     | disabled | disabled | disabled |
| l2cd-acc-aggr-cir-id-asc(vsa)  | n/c     | n/c     | disabled | disabled | disabled |
| l2cd-act-data-rate-up(vsa)     | n/c     | n/c     | disabled | disabled | disabled |
| l2cd-act-data-rate-dn(vsa)     | n/c     | n/c     | disabled | disabled | disabled |
| l2cd-min-data-rate-up(vsa)     | n/c     | n/c     | disabled | disabled | disabled |
| l2cd-min-data-rate-dn(vsa)     | n/c     | n/c     | disabled | disabled | disabled |
| l2cd-att-data-rate-up(vsa)     | n/c     | n/c     | disabled | disabled | disabled |
| l2cd-att-data-rate-dn(vsa)     | n/c     | n/c     | disabled | disabled | disabled |
| l2cd-max-data-rate-up(vsa)     | n/c     | n/c     | disabled | disabled | disabled |



|                                |         |         |          |          |          |
|--------------------------------|---------|---------|----------|----------|----------|
| l2cd-max-data-rate-dn(vsa)     | n/c     | n/c     | disabled | disabled | disabled |
| l2cd-min-lp-data-rate-up(vsa)  | n/c     | n/c     | disabled | disabled | disabled |
| l2cd-min-lp-data-rate-dn(vsa)  | n/c     | n/c     | disabled | disabled | disabled |
| l2cd-max-interlv-delay-up(vsa) | n/c     | n/c     | disabled | disabled | disabled |
| l2cd-act-interlv-delay-up(vsa) | n/c     | n/c     | disabled | disabled | disabled |
| l2cd-max-interlv-delay-dn(vsa) | n/c     | n/c     | disabled | disabled | disabled |
| l2cd-act-interlv-delay-dn(vsa) | n/c     | n/c     | disabled | disabled | disabled |
| l2cd-dsl-line-state(vsa)       | n/c     | n/c     | disabled | disabled | disabled |
| l2cd-dsl-type(vsa)             | n/c     | n/c     | disabled | disabled | disabled |
| l2tp-ppp-disconnect-cause      | n/c     | n/c     | n/c      | n/c      | disabled |
| mlppp-bundle-name              | n/c     | n/c     | enabled  | enabled  | enabled  |
| nas-identifier                 | enabled | enabled | enabled  | enabled  | enabled  |
| nas-port                       | n/c     | n/c     | enabled  | enabled  | enabled  |
| nas-port-id                    | n/c     | n/c     | enabled  | enabled  | enabled  |
| nas-port-type                  | n/c     | n/c     | enabled  | enabled  | enabled  |
| output-gigapkts(vsa)           | n/c     | n/c     | n/c      | n/c      | enabled  |
| output-gigawords               | n/c     | n/c     | n/c      | n/c      | enabled  |
| pcp-server-name(vsa)           | n/c     | n/c     | n/c      | disabled | disabled |
| pppoe-description(vsa)         | n/c     | n/c     | enabled  | enabled  | enabled  |
| profile-service-descr(vsa)     | n/c     | n/c     | disabled | disabled | disabled |
| qos-profile-name (vsa)         | n/c     | n/c     | n/c      | disabled | disabled |
| tunnel-assignment-id           | n/c     | n/c     | n/c      | enabled  | enabled  |
| tunnel-client-auth-id          | n/c     | n/c     | enabled  | enabled  | enabled  |
| tunnel-client-endpoint         | n/c     | n/c     | enabled  | enabled  | enabled  |
| tunnel-interface-id            | n/c     | n/c     | disabled | disabled | disabled |
| tunnel-medium-type             | n/c     | n/c     | enabled  | enabled  | enabled  |
| tunnel-preference              | n/c     | n/c     | n/c      | enabled  | enabled  |
| tunnel-server-attributes       | n/c     | n/c     | disabled | disabled | disabled |
| tunnel-server-auth-id          | n/c     | n/c     | enabled  | enabled  | enabled  |
| tunnel-server-endpoint         | n/c     | n/c     | enabled  | enabled  | enabled  |
| tunnel-type                    | n/c     | n/c     | enabled  | enabled  | enabled  |

**Meaning** Table 67 on page 297 lists the **show radius attributes-included** command output fields.

**Table 67: show radius attributes-included Output Fields**

| Field Name     | Field Description                                                                              |
|----------------|------------------------------------------------------------------------------------------------|
| Attribute Name | Name of the RADIUS attribute                                                                   |
| Account On     | Include status of the attribute in Acct-On messages: enabled, disabled, not configurable (n/c) |
| Account Off    | Include status of the attribute in Acct-Off messages: enabled, disabled, n/c                   |
| Access Request | Include status of the attribute in Access Request messages: enabled, disabled, n/c             |
| Account Start  | Include status of the attribute in Acct-Start messages: enabled, disabled, n/c                 |
| Account Stop   | Include status of the attribute in Acct-Stop messages: enabled, disabled, n/c                  |

- Related Documentation**
- [CLI Commands Used to Configure Juniper Networks VSAs on page 213](#)
  - [CLI Commands Used to Include ANCP-Related Juniper Networks VSAs in Access and Accounting Messages on page 215](#)
  - [CLI Commands Used to Include or Exclude Attributes in RADIUS Messages on page 218](#)
  - *show radius attributes-included*

---

## Monitoring Ignored RADIUS Attributes

---

**Purpose** Display the RADIUS attributes that are ignored in Access-Accept messages.

**Action** To display the RADIUS attributes that are ignored:

```
host1#show radius attributes-ignored
attribute framed-ip-netmask ignored from RADIUS server
attribute atm-service-category (vsa) accepted from RADIUS server
attribute atm-mbs (vsa) accepted from RADIUS server
attribute atm-pcr (vsa) accepted from RADIUS server
attribute atm-scr (vsa) accepted from RADIUS server
attribute egress-policy-name (vsa) accepted from RADIUS server
attribute ingress-policy-name (vsa) accepted from RADIUS server
attribute ipv6-egress-policy-name (vsa) accepted from RADIUS server
attribute ipv6-ingress-policy-name (vsa) accepted from RADIUS server
attribute virtual-router (vsa) accepted from RADIUS server
attribute pppoe-max-session (vsa) ignored from RADIUS server
```

- Related Documentation**
- [CLI Commands Used to Configure Juniper Networks VSAs on page 213](#)
  - [CLI Commands Used to Ignore Attributes when Receiving Access-Accept Messages on page 223](#)
  - *show radius attributes-ignored*

---

## Setting the Baseline for RADIUS Dynamic-Request Server Statistics

---

You can set a statistics baseline for packet mirroring-related RADIUS statistics. To show baseline statistics, use the **delta** keyword with the **show radius dynamic-request statistics** command.

To set a baseline for RADIUS statistics for packet mirroring:

- Issue the **baseline radius dynamic-request** command:

```
host1#baseline radius dynamic-request
```

There is no **no** version.

- Related Documentation**
- [Monitoring RADIUS Dynamic-Request Server Statistics on page 299](#)
  - *baseline radius dynamic-request*

## Monitoring RADIUS Dynamic-Request Server Statistics

**Purpose** Display RADIUS dynamic-request server statistics.

**Action** To display RADIUS dynamic-request statistics:

```
host1#show radius dynamic-request statistics
```

```

RADIUS Request Statistics

Statistic 10.10.3.4

UDP Port 1700
Disconnect Requests 0
Disconnect Accepts 0
Disconnect Rejects 0
Disconnect No Session ID 0
Disconnect Bad Authenticators 0
Disconnect Packets Dropped 0
COA Requests 0
COA Accepts 0
COA Rejects 0
COA No Session ID 0
COA Bad Authenticators 0
COA Packets Dropped 0
No Secret 0
Unknown Request 0

Invalid Addresses Received :0

```

**Meaning** [Table 68 on page 299](#) lists the `show radius dynamic-request statistics` command output fields.

**Table 68: show radius dynamic-request statistics Output Fields**

| Field Name                           | Field Description                                                                                                                                 |
|--------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| Udp Port                             | Port on which the router listens for RADIUS server                                                                                                |
| Disconnect or COA Requests           | RADIUS-initiated disconnect or COA requests received                                                                                              |
| Disconnect or COA Accepts            | RADIUS-initiated disconnect or COA requests accepted                                                                                              |
| Disconnect or COA Rejects            | RADIUS-initiated disconnect or COA requests rejected                                                                                              |
| Disconnect or COA No Session ID      | RADIUS-initiated disconnect or COA messages rejected because the request did not include a session ID attribute                                   |
| Disconnect or COA Bad Authenticators | RADIUS-initiated disconnect or COA messages rejected because the calculated authenticator in the authenticator field of the request did not match |
| Disconnect or COA Packets Dropped    | RADIUS-initiated disconnect or COA packets dropped because of queue overflow                                                                      |

**Table 68: show radius dynamic-request statistics Output Fields (continued)**

| Field Name                 | Field Description                                                                             |
|----------------------------|-----------------------------------------------------------------------------------------------|
| No Secret                  | Messages rejected because a secret was not present in the authenticator field                 |
| Unknown Requests           | Packets received with an invalid RADIUS code for RADIUS disconnect or change of authorization |
| Invalid Addresses Received | Number of invalid addresses received                                                          |

- Related Documentation**
- [Setting the Baseline for RADIUS Dynamic-Request Server Statistics on page 298](#)
  - *show radius statistics*

## Monitoring the Configuration of the RADIUS Dynamic-Request Server

**Purpose** Display the configuration of the RADIUS dynamic-request server.

**Action** To display the configuration of the RADIUS dynamic-request server:

```
host1#show radius dynamic-request servers
```

```

RADIUS Request Configuration

 IP Address Udp Port Disconnect Change Of Secret
 ----- -
192.168.2.3 1700 disabled disabled <NULL>
10.10.120.104 1700 disabled disabled mysecret

```

**Meaning** [Table 69 on page 300](#) lists the **show radius dynamic-request servers** command output fields.

**Table 69: show radius dynamic-request servers Output Fields**

| Field Name              | Field Description                                  |
|-------------------------|----------------------------------------------------|
| IP address              | IP address of the RADIUS server                    |
| Udp Port                | Port on which the router listens for RADIUS server |
| Disconnect              | Status of RADIUS-initiated disconnect feature      |
| Change of Authorization | Status of change of authorization feature          |
| Secret                  | Secret used to connect to RADIUS server            |

**Related Documentation**

- *show radius servers*

## Setting a Baseline for RADIUS Relay Statistics

You can set a baseline for RADIUS relay statistics. To show baseline statistics, use the **delta** keyword with the **show radius relay** command.

To set a baseline for RADIUS relay statistics:

- Issue the **baseline radius relay** command:

```
host1#baseline radius relay
```

There is no **no** version.

**Related Documentation**

- [Monitoring RADIUS Relay Server Statistics on page 301](#)
- *baseline radius relay*

## Monitoring RADIUS Relay Server Statistics

**Purpose** Display RADIUS relay server statistics.

**Action** To show RADIUS relay server statistics that were baselined:

```
host1#show radius relay statistics delta
```

```
RADIUS Relay Authentication Server Statistics
```

```

Statistic Total

Access Requests 1000
Access Accepts 1000
Access Challenges 0
Access Rejects 0
Pending Requests 0
Duplicate Requests 0
Malformed Requests 0
Bad Authenticators 0
Unknown Requests 0
Dropped Packets 0
Invalid Requests 0
```

```
Statistics baseline set FRI APR 02 2004 19:01:52 UTC
```

```
RADIUS Relay Accounting Server Statistics
```

```

Statistic Total

Accounting Requests 1000
Start 1000
Stop 0
Interim 0
Accounting Responses 1000
Start 1000
Stop 0
Interim 0
Pending Requests 0
```

```

Duplicate Requests 0
Malformed Requests 0
Bad Authenticators 0
Unknown Requests 0
Dropped Packets 0
Invalid Requests 0
Statistics baseline set FRI APR 02 2004 19:01:52 UTC

```

**Meaning** [Table 70 on page 302](#) lists the **show radius relay statistics** command output fields.

**Table 70: show radius relay statistics Output Fields**

| Field Name           | Field Description                                                                                                                                   |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| Access Requests      | Number of access requests received                                                                                                                  |
| Access Accepts       | Number of access accepts received                                                                                                                   |
| Access Challenges    | Number of access challenges received                                                                                                                |
| Access Rejects       | Number of access rejects received                                                                                                                   |
| Pending Requests     | Number of access requests waiting for a response                                                                                                    |
| Duplicate Requests   | Number of duplicate requests received while the previous request is pending                                                                         |
| Malformed Requests   | Requests with attributes having an invalid length or unexpected attributes                                                                          |
| Bad Authenticators   | Authenticator in the response is incorrect for the matching request; can occur if the secret for the RADIUS relay server and the WAP does not match |
| Unknown Requests     | Packets received from nonconfigured clients                                                                                                         |
| Dropped Packets      | Packets dropped because of queue overflow                                                                                                           |
| Invalid Requests     | Number of invalid requests received                                                                                                                 |
| Accounting Requests  | Number of accounting requests received, broken down by type of request                                                                              |
| Accounting Responses | Number of accounting responses, broken down by type of request                                                                                      |

**Related Documentation**

- [Setting a Baseline for RADIUS Relay Statistics on page 301](#)
- *show radius relay statistics*

## Monitoring the Configuration of the RADIUS Relay Server

**Purpose** Display information about the RADIUS relay server configuration.

**Action** To display the RADIUS relay server configuration:

```
host1#show radius relay servers
```

```
RADIUS Relay Authentication Server Configuration
```

```

IP Address IP Mask Secret

10.10.8.15 255.255.255.255 newsecret
192.168.102.5 255.255.255.255 999Y2K
Udp Port: 1812
```

```
RADIUS Relay Accounting Server Configuration
```

```

IP Address IP Mask Secret

10.10.1.0 255.255.255.0 N08pxq
192.168.102.5 255.255.255.255 12BE$56
Udp Port: 1813
```

**Meaning** [Table 71 on page 303](#) lists the **show radius relay servers** command output fields.

**Table 71: show radius relay servers Output Fields**

| Field Name | Field Description                                                    |
|------------|----------------------------------------------------------------------|
| IP Address | Address of the RADIUS relay server                                   |
| IP Mask    | Mask of the RADIUS relay server                                      |
| Secret     | Secret used for exchanges between the RADIUS relay server and client |
| Udp Port   | Router's port on which the RADIUS relay server listens               |

**Related Documentation**

- *show radius relay servers*

## Monitoring the Status of RADIUS Relay UDP Checksums

**Purpose** Display status of RADIUS relay UDP checksums.

**Action** To display the status of UDP checksums:

```
host1(config)#show radius relay udp-checksum
udp-checksums enabled
```

**Meaning** [Table 72 on page 304](#) lists the **show radius relay udp-checksum** command output fields.

Table 72: show radius relay udp-checksum Output Fields

| Field Name    | Field Description                            |
|---------------|----------------------------------------------|
| udp-checksums | Status of UDP checksums: enabled or disabled |

**Related Documentation**     • *show radius relay udp-checksum*

---

## Monitoring the Status of ICR Partition Accounting

---

**Purpose**     Display the status of ICR partition accounting.

**Action**     To display the status of ICR partition accounting:  
  
host1#**show radius icr-partition-accounting**  
enabled

**Meaning**     ICR partition accounting status is either enabled or disabled.

**Related Documentation**     • *show radius icr-partition-accounting*



## CHAPTER 10

# Configuring TACACS+

This chapter explains how to enable and configure TACACS+ in your E Series router. It has the following sections:

- [Understanding TACACS+ on page 305](#)
- [TACACS+ Platform Considerations on page 309](#)
- [TACACS+ References on page 309](#)
- [Retry Attempts for Successful TCP Connection Overview on page 310](#)
- [Configuring TACACS+ on page 310](#)

## Understanding TACACS+

---

With the increased use of remote access, the need for managing more network access servers (NAS) has increased. Additionally, the need for control access on a per-user basis has escalated, as has the need for central administration of users and passwords.

Terminal Access Controller Access Control System (TACACS) is a security protocol that provides centralized validation of users who are attempting to gain access to a router or NAS. TACACS+, a more recent version of the original TACACS protocol, provides separate authentication, authorization, and accounting (AAA) services.



**NOTE:** TACACS+ is a completely new protocol and is not compatible with TACACS or XTACACS.

The TACACS+ protocol provides detailed accounting information and flexible administrative control over the authentication, authorization, and accounting process. The protocol allows a TACACS+ client to request detailed access control and allows the TACACS+ process to respond to each component of that request. TACACS+ uses Transmission Control Protocol (TCP) for its transport.

TACACS+ provides security by encrypting all traffic between the NAS and the process. Encryption relies on a secret key that is known to both the client and the TACACS+ process.

[Table 73 on page 306](#) describes terms that are frequently used in this chapter.

Table 73: TACACS-Related Terms

| Term            | Description                                                                                                                                                                                                                                                                                                                                      |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| NAS             | Network access server. A device that provides connections to a single user, to a network or subnetwork, and to interconnected networks. In reference to TACACS+, the NAS is the E Series router.                                                                                                                                                 |
| TACACS+ process | A program or software running on a security server that provides AAA services using the TACACS+ protocol. The program processes authentication, authorization, and accounting requests from an NAS. When processing authentication requests, the process might respond to the NAS with a request for additional information, such as a password. |
| TACACS+ host    | The security server on which the TACACS+ process is running. Also referred to as a TACACS+ server.                                                                                                                                                                                                                                               |

## AAA Overview

TACACS+ allows effective communication of AAA information between NASs and a central server. The separation of the AAA functions is a fundamental feature of the TACACS+ design:

- **Authentication**—Determines who a user is, then determines whether that user should be granted access to the network. The primary purpose is to prevent intruders from entering your networks. Authentication uses a database of users and passwords.
- **Authorization**—Determines what an authenticated user is allowed to do. Authorization gives the network manager the ability to limit network services to different users. Also, the network manager can limit the use of certain commands to various users. Authorization cannot occur without authentication.
- **Accounting**—Tracks what a user did and when it was done. Accounting can be used for an audit trail or for billing for connection time or resources used. Accounting can occur independent of authentication and authorization.

Central management of AAA means that the information is in a single, centralized, secure database, which is much easier to administer than information distributed across numerous devices. Both RADIUS and TACACS+ protocols are client-server systems that allow effective communication of AAA information.

For information about RADIUS, see [“RADIUS Overview” on page 180](#).

## Administrative Login Authentication

Fundamentally, TACACS+ provides the same services as RADIUS. Every authentication login attempt on an NAS is verified by a remote TACACS+ process.

TACACS+ authentication uses three packet types. Start packets and Continue packets are always sent by the user. Reply packets are always sent by the TACACS+ process.

TACACS+ sets up a TCP connection to the TACACS+ host and sends a Start packet. The TACACS+ host responds with a Reply packet, which either grants or denies access, reports an error, or challenges the user.

TACACS+ might challenge the user to provide username, password, passcode, or other information. Once the requested information is entered, TACACS+ sends a Continue packet over the existing connection. The TACACS+ host sends a Reply packet. Once the authentication is complete, the connection is closed. Only three login retries are allowed.

To enable login authentication through both TACACS+ and RADIUS servers, use the **aaa new-model** command to specify AAA authentication for Telnet sessions.

## Privilege Authentication

The privilege authentication process determines whether a user is allowed to use commands at a particular privilege level. This authentication process is handled similarly to login authentication, except that the user is limited to one authentication attempt. An empty reply to the challenge forces an immediate access denial. The **aaa authentication enable default** command allows you to set privilege authentication for users.

## Login Authorization

To allow login authorization through the TACACS+ server, you can use the following commands: **aaa authorization**, **aaa authorization config-commands**, and **authorization**. For information about using these commands, see the *Passwords and Security* chapter in the *JunosE System Basics Configuration Guide*.

## Accounting

The TACACS+ accounting service enables you to create an audit trail of User Exec sessions and command-line interface (CLI) commands that have been executed within these sessions. For example, you can track user CLI connects and disconnects, when configuration modes have been entered and exited, and which configuration and operational commands have been executed.

You configure TACACS+ accounting in the JunosE Software by defining accounting method lists and then associating consoles and lines with the method lists. You define an accounting method list with a service type, name, accounting mode, and method:

- service type—Specifies the type of information being recorded
- name—Uniquely identifies an accounting method list within a service type
- accounting mode—Specifies what type of accounting records will be generated
- method—Specifies the protocol for sending the accounting records to a security server

You can then configure consoles and lines with an accounting method list name for each service type:

- Method list—A specified configuration that defines how the NAS performs the AAA accounting service. A service type can be configured with multiple method lists with different names, and a method list name can be used for different service types. Initially, no accounting method list is defined; therefore TACACS+ accounting is disabled.
- Default method list—Configuration used by consoles and lines when no named method list is assigned. You enable TACACS+ accounting by defining default accounting method lists for each service type.

- Named method list—Assigned to a console, specific line, or group of lines; overrides the default method list.
- Service type—Specifies the type of information provided by the TACACS+ accounting service:
  - Exec—Provides information about User Exec terminal sessions, such as telnet, Local Area Transport (LAT), and rlogin, on the NAS.
  - Commands <0-15>—Provides information about User Exec mode CLI commands for a specified privilege level that are being executed on the NAS. Each of the sixteen command privilege levels is a separate service type. Accounting records are generated for commands executed by users, CLI scripts, and macros.
- Accounting mode—Specifies the type of accounting records that are recorded on the TACACS+ server. Accounting records track user actions and resource usage. You can analyze and use the records for network management, billing, and auditing purposes.
  - start-stop—A start accounting record is generated just before a process begins, and a stop accounting record is generated after a process successfully completes. This mode is supported only for the Exec service type.
  - stop-only—A stop accounting record is generated after a process successfully completes. This mode is supported only for the Commands service types.

The NAS sends TACACS+ accounting packets to the TACACS+ host. The accounting packets contain data in the packet header, packet body, and attribute-value pairs (AVPs). [Table 74 on page 308](#) provides descriptions of the TACACS+ accounting data.

**Table 74: TACACS+ Accounting Information**

| Field/Attribute | Location      | Description                                                                                     |
|-----------------|---------------|-------------------------------------------------------------------------------------------------|
| major_version   | Packet header | Major TACACS+ version number                                                                    |
| minor_version   | Packet header | Minor TACACS+ version number                                                                    |
| type            | Packet header | Type of the AAA service: Accounting                                                             |
| flags           | Packet body   | Bitmapped flags representing the record type: start accounting record or stop accounting record |
| priv-level      | Packet body   | Privilege level of the user executing the Exec session or CLI command: 0 - 15                   |
| user            | Packet body   | Name of user running the Exec session or CLI command                                            |
| port            | Packet body   | NAS port used by the Exec session or CLI command                                                |
| rem-addr        | Packet body   | User's remote location; either an IP address or the caller ID                                   |
| service         | AVP           | User's primary service: Shell                                                                   |

Table 74: TACACS+ Accounting Information (*continued*)

| Field/Attribute | Location | Description                                                                                                                  |
|-----------------|----------|------------------------------------------------------------------------------------------------------------------------------|
| cmd             | AVP      | CLI command that is to be executed: specified for Command-level accounting only                                              |
| task_id         | AVP      | Unique sequential identifier used to match start and stop records for a task                                                 |
| elapsed_time    | AVP      | Elapsed time in seconds for the task execution: specified for Exec-level accounting stop records only                        |
| timezone        | AVP      | Time zone abbreviation used “ <a href="#">Monitoring TACACS+ Statistics</a> ” on <a href="#">page 315</a> for all timestamps |

**Related  
Documentation**

- [Configuring TACACS+ on page 310](#)
- [Monitoring TACACS+ Statistics on page 315](#)
- [Monitoring TACACS+ Information on page 317](#)

## TACACS+ Platform Considerations

TACACS+ is supported on all E Series routers. For information about the modules supported on E Series routers:

- See the *ERX Module Guide* for modules supported on ERX7xx models, ERX14xx models, and the ERX310 Broadband Services Router.
- See the *E120 and E320 Module Guide* for modules supported on the E120 and E320 Broadband Services Routers.

**Related  
Documentation**

- [Understanding TACACS+ on page 305](#)

## TACACS+ References

For additional information about the TACACS+ protocol, see the following resources:

- The TACACS+ Protocol, Version 1.78—draft-grant-tacacs-02.txt (January 1997 expiration)
- RFC 2865—Remote Authentication Dial In User Service (RADIUS) (June 2000)



**NOTE:** IETF drafts are valid for only 6 months from the date of issuance. They must be considered as works in progress. Please refer to the IETF website at <http://www.ietf.org> for the latest drafts.

- Related Documentation**
- [Understanding TACACS+ on page 305](#)

---

## Retry Attempts for Successful TCP Connection Overview

The TACACS+ client uses Transmission Control Protocol (TCP) to establish a connection with the TACACS+ server. TCP provides a timeout value of 75 seconds. Within the stipulated time, the TCP connection between the TACACS+ client and the TACACS+ server may not be established due to certain conditions, such as network congestion or connectivity failure. If the TACACS+ client fails to establish a TCP connection, the TACACS+ commands are not executed.

For robust TACACS+ operations, the TACACS+ client can reattempt connection upon connection failure. You can configure the TACACS+ client retry value to establish the TCP connection with the TACACS+ server. The retry is initiated in either of the following two cases:

- User configuration timeout occurs
- TCP connection timeout occurs



**NOTE:** The same TACACS+ server is reattempted if TCP connection is not established or the user-configured timer expires.

---

The maximum retry attempt for a request is five. By default, the maximum retry value is two.

- Related Documentation**
- [Configuring TACACS+ on page 310](#)
  - *tacacs-server retransmit-retries*

---

## Configuring TACACS+

Terminal Access Controller Access Control System (TACACS) is a security protocol that provides centralized validation of users who are attempting to gain access to a router or NAS. TACACS+, a more recent version of the original TACACS protocol, provides separate authentication, authorization, and accounting (AAA) services. This topic includes the following tasks:

1. [Configuring TACACS+ Support on page 310](#)
2. [Configuring Authentication on page 311](#)
3. [Configuring Accounting on page 312](#)

### Configuring TACACS+ Support

Before you begin to configure TACACS+, you must determine the following for the TACACS+ authentication and accounting servers:

- IP addresses

- TCP port numbers
- Secret keys

To use TACACS+, you must enable AAA. To configure your router to support TACACS+, perform the following tasks. Some of the tasks are optional. Once you configure TACACS+ support on the router, you can configure TACACS+ authentication, authorization, and accounting independent of each other.

You can configure the TACACS+ server only on default virtual routers. If you attempt to configure TACACS+ server settings on VRs other than the default VR or in a VRF, an error message is displayed.

1. Specify the names of the IP host or hosts maintaining a TACACS+ server. Optionally, you can specify other parameters, such as port number, timeout interval, and key.

```
host1(config)#tacacs-server host 192.168.1.27 port 10 timeout 3 key your_secret primary
```

2. (Optional) Set the authentication and encryption key value shared by all TACACS+ servers that do not have a server-specific key set up by the **tacacs-server host** command.

```
host1(config)#tacacs-server key "͹P^"
```

3. (Optional) Set alternative source addresses to be used for TACACS+ server communications.

```
host1(config)#tacacs-server source-address 192.168.134.63
```

4. (Optional) Set the timeout value for all TACACS+ servers that do not have a server-specific timeout set up by the **tacacs-server host** command.

```
host1(config)#tacacs-server timeout 15
```

5. (Optional) Set the retry value for a TACACS+ client. The maximum retry attempt for a request is five. By default, the retry value is two.

```
host1(config)#tacacs-server retransmit-retries 4
```

## Configuring Authentication

Once TACACS+ support is enabled on the router, you can configure TACACS+ authentication. Perform the following steps:

1. Specify AAA new model as the authentication method for the vty lines on your router.

```
host1(config)#aaa new-model
```

2. Specify AAA authentication by defining an authorization methods list.

```
host1(config)#aaa authentication login tac tacacs+ radius enable
```

3. Specify the privilege level by defining a methods list that uses TACACS+ for authentication.

```
host1(config)#aaa authentication enable default tacacs+ radius enable
```

4. Configure vty lines.

```
host1(config)#line vty 0 4
```

5. Apply an authentication list to the vty lines you specified on your router.

```
host1(config-line)#login authentication tac
```

## Configuring Accounting

Once TACACS+ support is enabled on the router, you can configure TACACS+ accounting. Perform the following steps:

1. Specify AAA new model as the accounting method for your router.

```
host1(config)#aaa new-model
```

2. Enable TACACS+ accounting on the router, and configure accounting method lists. For example:

```
host1(config)#aaa accounting exec default start-stop tacacs+
host1(config)#aaa accounting commands 0 listX stop-only tacacs+
host1(config)#aaa accounting commands 1 listX stop-only tacacs+
host1(config)#aaa accounting commands 13 listY stop-only tacacs+
host1(config)#aaa accounting commands 14 default stop-only tacacs+
host1(config)#aaa accounting commands 15 default stop-only tacacs+
```

3. (Optional) Specify that accounting records are not generated for users without explicit usernames.

```
host1(config)#aaa accounting suppress null-username
```

4. Apply accounting method lists to a console or lines. For example:

```
host1(config)#line console 0
host1(config-line)#accounting commands 0 listX
host1(config-line)#accounting commands 1 listX
host1(config-line)#accounting commands 13 listY
host1(config-line)#exit
host1(config)#line vty 0 4
host1(config-line)#accounting commands 13 listY
```

Note that Exec accounting and User Exec mode commands accounting for privilege levels 14 and 15 are now enabled for all lines and consoles with the creation of their default method list, as shown in Step 2.

### Related Documentation

- *aaa accounting commands*
- *aaa accounting exec*
- *aaa accounting suppress null-username*
- *aaa authentication enable default*
- *aaa authentication login*
- *aaa new-model*
- *line*
- *login authentication*
- *tacacs-server host*



- *tacacs-server key*
- *tacacs-server source-address*
- *tacacs-server timeout*



## CHAPTER 11

# Monitoring TACACS+

This chapter describes how to monitor the current TACACS+ configurations.

TACACS+ topics are described in the following sections:

- [Setting Baseline TACACS+ Statistics on page 315](#)
- [Monitoring TACACS+ Statistics on page 315](#)
- [Monitoring TACACS+ Information on page 317](#)

## Setting Baseline TACACS+ Statistics

---

You can set a baseline for TACACS+ statistics.

To set the baseline:

- Issue the **baseline tacacs** command:

```
host1#baseline tacacs
```

There is no **no** version.

**Related Documentation**

- [baseline tacacs](#)

## Monitoring TACACS+ Statistics

---

**Purpose** Display TACACS+ statistics.

**Action** To display TACACS+ statistics:

```
host1#show statistics tacacs
```

```
TACACSPPLUS Statistics

Statistic 10.5.0.174 10.5.1.199

Search Order 1 2
TCP Port 3049 4049
Auth Requests 140 0
Auth Replies 85 0
Auth Pending 43 0
Auth Timeouts 12 0
Author Requests 6399 97
Author Replies 6301 0
```

|                 |      |    |
|-----------------|------|----|
| Author Pending  | 0    | 0  |
| Author Timeouts | 98   | 97 |
| Acct Requests   | 6321 | 37 |
| Acct Replies    | 6280 | 0  |
| Acct Pending    | 4    | 0  |
| Acct Timeouts   | 37   | 37 |

**Meaning** [Table 75 on page 316](#) lists the **show statistics tacacs** command output fields.

**Table 75: show statistics tacacs Output Fields**

| Field Name      | Field Description                                                          |
|-----------------|----------------------------------------------------------------------------|
| Statistic       | IP address of the host                                                     |
| Search Order    | The order in which requests are sent to hosts until a response is received |
| TCP Port        | TCP port of the host                                                       |
| Auth Requests   | Number of authentication requests sent to the host                         |
| Auth Replies    | Number of authentication replies received from the host                    |
| Auth Pending    | Number of expected but not received authentication replies from the host   |
| Auth Timeouts   | Number of authentication timeouts for the host                             |
| Author Requests | Number of authorization requests sent to the host                          |
| Author Replies  | Number of authorization replies received from the host                     |
| Author Pending  | Number of expected but not received authorization replies from the host    |
| Author Timeouts | Number of authorization timeouts for the host                              |
| Acct Requests   | Number of accounting requests sent to the host                             |
| Acct Replies    | Number of accounting replies received from the host                        |
| Acct Pending    | Number of expected but not received accounting replies from the host       |
| Acct Timeouts   | Number of accounting timeouts for the host                                 |

**Related Documentation**

- *show statistics tacacs*

## Monitoring TACACS+ Information

**Purpose** Display TACACS+ information.

**Action** To display TACACS+ information.

```
host1#show tacacs
Key = hippo
Timeout = <NOTSET>, built-in timeout of 5 will be used
Source-address = <NOTSET>
Retry-attempts = 3
```

TACACS+ Configuration, (\*) denotes inherited

| IP Address | Tcp Port | Timeout | Primary | Key       | Search Order |
|------------|----------|---------|---------|-----------|--------------|
| 10.5.0.174 | 3049     | 5 (*)   | y       | hippo (*) | 1            |
| 10.5.1.199 | 1049     | 5 (*)   | n       | hippo (*) | 2            |

To display overall statistics:

```
host1#show tacacs statistics
```

To display statistics since they were baselined; deltas are not calculated for the pending statistics:

```
host1#show tacacs delta
```

**Meaning** [Table 76 on page 317](#) lists the **show tacacs** command output fields.

**Table 76: show tacacs Output Fields**

| Field Name               | Field Description                                                                                                        |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------|
| Key                      | Authentication and encryption key                                                                                        |
| Timeout                  | TACACS+ host response timeout in seconds                                                                                 |
| Source-address           | Alternative source IP address configured                                                                                 |
| Retry-attempts           | Number of retry attempts that will be made to establish a TCP connection between a TACACS+ client and the TACACS+ server |
| TACACSPLUS Configuration | Table contains statistics for each host                                                                                  |
| IP Address               | IP address of the host                                                                                                   |
| TCP Port                 | TCP port of the host for each IP address                                                                                 |
| Timeout                  | Timeout interval in seconds for each IP address                                                                          |
| Primary                  | This IP address's primary host; options: y = yes, n = no                                                                 |

**Table 76: show tacacs Output Fields (*continued*)**

| Field Name   | Field Description                                                          |
|--------------|----------------------------------------------------------------------------|
| Key          | Authentication and encryption key for this IP address                      |
| Search Order | The order in which requests are sent to hosts until a response is received |

**Related Documentation**

- *show tacacs*

## PART 3

# Managing L2TP

- [L2TP Overview on page 321](#)
- [Configuring an L2TP LAC on page 331](#)
- [Configuring an L2TP LNS on page 365](#)
- [Configuring L2TP Dial-Out on page 407](#)
- [L2TP Disconnect Cause Codes on page 421](#)
- [Monitoring L2TP and L2TP Dial-Out on page 427](#)





## CHAPTER 12

# L2TP Overview

Layer 2 Tunneling Protocol (L2TP) is a client-server protocol that allows Point-to-Point Protocol (PPP) to be tunneled across a network. This chapter includes the following topics that provide information for configuring L2TP on the Juniper Networks E Series Broadband Services Routers.

- [L2TP Overview on page 321](#)
- [L2TP Terminology on page 322](#)
- [Implementing L2TP on page 323](#)
- [Packet Fragmentation on page 325](#)
- [L2TP Platform Considerations on page 326](#)
- [L2TP Module Requirements on page 326](#)
- [Sessions and Tunnels Supported on page 328](#)
- [L2TP References on page 329](#)

## L2TP Overview

---

L2TP encapsulates layer 2 packets, such as PPP, for transmission across a network. An L2TP access concentrator (LAC), configured on an access device, such as an E Series router, receives packets from a remote client and forwards them to an L2TP network server (LNS), on a remote network.

You can configure your router to act as an LAC in pass-through mode in which the LAC receives packets from a remote client and then forwards them at layer 2 directly to the LNS.

The E Series router creates tunnels dynamically by using authentication, authorization, and accounting (AAA) authentication parameters and transmits L2TP packets to the LNS via IP/User Datagram Protocol (UDP). Traffic travels in an L2TP *session*. A tunnel is an aggregation of one or more sessions. [Figure 7 on page 322](#) and [Figure 8 on page 322](#) show the E Series router in typical LAC and LNS arrangements.

Figure 7: Using the E Series Router as an LAC





Figure 8: Using the E Series Router as an LNS



**NOTE:** The E Series router does not support terminating both ends of a tunnel or session in the same router.

#### Related Documentation

- [Implementing L2TP on page 323](#)
- [L2TP Platform Considerations on page 326](#)
- [L2TP References on page 329](#)

## L2TP Terminology

Table 77 on page 322 describes the basic terms for L2TP.

Table 77: L2TP Terms

| Term                       | Description                                                                                                                                                                                            |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Attribute value pair (AVP) | Combination of a unique attribute—represented by an integer—and a value containing the actual value identified by the attribute.                                                                       |
| LAC                        | L2TP access concentrator (LAC)—a node that acts as one side of an L2TP tunnel endpoint and is a peer to the LNS. An LAC sits between an LNS and a remote system and forwards packets to and from each. |
| Call                       | A connection (or attempted connection) between a remote system and an LAC.                                                                                                                             |

Table 77: L2TP Terms (*continued*)

| Term                 | Description                                                                                                                                                                                                                                                                   |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| LNS                  | L2TP network server (LNS)—a node that acts as one side of an L2TP tunnel endpoint and is a peer to the LAC. An LNS is the logical termination point of a PPP connection that is being tunneled from the remote system by the LAC.                                             |
| Peer                 | In the L2TP context, refers to either the LAC or LNS. An LAC's peer is an LNS, and vice versa.                                                                                                                                                                                |
| Proxy authentication | Authentication data from the PPP client that is sent from the LNS as part of a proxy LCP. Data might include attributes such as authentication type, authentication name, and authentication challenge.                                                                       |
| Proxy LCP            | LCP (Link Control Protocol) negotiation that is performed by the LAC on behalf of the LNS. Proxy sent by the LAC to the LNS containing attributes such as the last configuration attributes sent and received from the client.                                                |
| Remote system        | An end-system or router attached to a remote access network, which is either the initiator or recipient of a call.                                                                                                                                                            |
| Session              | A logical connection created between the LAC and the LNS when an end-to-end PPP connection is established between a remote system and the LNS.<br><br><b>NOTE:</b> There is a one-to-one relationship between established L2TP sessions and their associated PPP connections. |
| Tunnel               | A connection between an LAC-LNS pair consisting of a control connection and 0 or more L2TP sessions.                                                                                                                                                                          |

**Related Documentation**

- [L2TP Overview on page 321](#)

## Implementing L2TP

The implementation of L2TP for the E Series router uses four levels:

- System—The router
- Destination—The remote L2TP system
- Tunnel—A direct path between the LAC and the LNS
- Session—A PPP connection in a tunnel

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

## Sequence of Events on the LAC

The E Series router creates destinations, tunnels, and sessions dynamically, as follows:

1. The client initiates a PPP connection with the router.
2. The router and the client exchange Link Control Protocol (LCP) packets. For details about negotiating PPP connections, see the *Configuring Point-to-Point Protocol* chapter in the *JunosE Link Layer Configuration Guide*.
3. By using either a local database related to the domain name or RADIUS authentication, the router determines either to terminate or to tunnel the PPP connection.
4. If the router discovers that it should tunnel the session, it does the following:
  - a. Sets up a new destination or selects an existing destination.
  - b. Sets up a new tunnel or selects an existing tunnel.
  - c. Opens a new session.
5. The router forwards the results of the LCP negotiations and authentication to the LNS.

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



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

---

## Sequence of Events on the LNS

The E Series router sets up an LNS as follows:

1. An LAC initiates a tunnel with the router.
2. The router verifies that a tunnel with this LAC is valid—destination configured, hostname and tunnel password correct.
3. The router completes the tunnel setup with the LAC.
4. The LAC sets up a session with the router.
5. The router creates a dynamic PPP interface on top of the session.
6. If they are enabled and present, the router takes the proxy LCP and the proxy authentication data and passes them to PPP.
7. The E Series PPP processes the proxy LCP, if it is present, and, if acceptable, places LCP on the router in opened state without renegotiation of LCP.



**NOTE:** If proxy LCP is not present or not acceptable, the router negotiates LCP with the remote system.

8. The E Series PPP processes the proxy authentication data, if it is present, and passes the data to AAA for verification. (If the data is not present, E Series PPP requests the data from the remote system.)
9. The router passes the authentication results to the remote system.

**Related  
Documentation**

- [L2TP Overview on page 321](#)
- [L2TP Platform Considerations on page 326](#)
- [L2TP Module Requirements on page 326](#)
- [L2TP References on page 329](#)

## Packet Fragmentation

The E Series router supports the reassembly of IP-fragmented L2TP packets. (For more information, see the *IP Reassembly for Tunnels* chapter in the *JunosE IP Services Configuration Guide*.) However, it is preferable to prevent fragmentation within L2TP tunnels because of the effects of fragmentation and reassembly on performance.

To prevent fragmentation, PPP LCP negotiation of the maximum receive unit (MRU) may be used to determine a proper maximum transmission unit (MTU). However, the normal automatic method of determining the proper MRU to negotiate (by evaluating the MRU of all lower layers in the interface stack) is not adequate for L2TP. The initial LCP negotiation between PPP in the client and the LAC is inadequate because it does not cover the entire extent of the eventual PPP session that travels all the way from the client to the LNS. Furthermore, even if PPP in the LNS chooses to renegotiate the MRU, it has no way to determine the proper MRU, since it does not know the minimum MRU on all of the intervening links between it and the LAC.

To overcome the inadequacy of normal determination of the MRU under such circumstances, you can configure the PPP MRU size by using the **ppp mru** command in Profile Configuration mode, Interface Configuration mode, or Subinterface Configuration mode. Use Profile Configuration mode for dynamic PPP interfaces, and Interface Configuration mode or Subinterface Configuration mode for static PPP interfaces.

When you specify the size, you need to take into account the MRU for all possible links between the LAC and the LNS. You must also take into account the L2TP encapsulation that is added to all packets entering the tunnel.

For example, if the link between the LAC and LNS with the lowest MRU were an Ethernet link, the following calculation applies:

|                                                                   |      |
|-------------------------------------------------------------------|------|
| Minimum link MRU                                                  | 1500 |
| L2TP encapsulating IP header                                      | -20  |
| L2TP encapsulating UDP header                                     | -8   |
| Maximum L2TP header (assumes a maximum of 16 bytes of Offset Pad) | -30  |

|                     |      |
|---------------------|------|
| MRU size to specify | 1442 |
|---------------------|------|

If the smallest intervening link is an Ethernet link, specifying **ppp mru 1442** at either the LAC or LNS guarantees that no fragmentation will occur within the L2TP tunnel.

- Related Documentation**
- [L2TP Overview on page 321](#)
  - [Implementing L2TP on page 323](#)

---

## L2TP Platform Considerations

For information about modules that support LNS and LAC on the ERX7xx models, ERX14xx models, and the ERX310 Broadband Services Router:

- See *ERX Module Guide, Table 1, ERX Module Combinations* for detailed module specifications.
- See *ERX Module Guide, Appendix A, Module Protocol Support* for information about the modules that support LNS and LAC.

For information about modules that support LNS and LAC on the E120 and E320 Broadband Services Routers:

- See *E120 and E320 Module Guide, Table 1, Modules and IOAs* for detailed module specifications.
- See *E120 and E320 Module Guide, Appendix A, IOA Protocol Support* for information about the modules that support LNS and LAC.

- Related Documentation**
- [L2TP Overview on page 321](#)
  - [L2TP References on page 329](#)

---

## L2TP Module Requirements

The supported modules for LNS depends on the type of E Series router that you have.

### ERX7xx Models, ERX14xx Models, and the ERX310 Router

To use an LNS on ERX7xx models, ERX14xx models, and the ERX310 router, at least one Service line module (SM) or a module that supports the use of shared tunnel-server ports must be installed in the ERX router. For information about installing modules in the ERX router, see the *ERX Hardware Guide*.

SMs provide dedicated tunnel-server ports that are always configured on the module. Unlike other line modules, SMs do not pair with corresponding I/O modules that contain ingress and egress ports. Instead, they receive data from and transmit data to other line modules with access to ingress and egress ports on their own associated I/O modules.

You can also create tunnels on E Series modules that support shared tunnel-server ports. You can configure (provision) a shared tunnel-server port to use a portion of the module's

bandwidth to provide tunnel services. For a list of the modules that support shared tunnel-server ports, see the *ERX Module Guide*.

When you configure the GE-2 line module or the GE-HDE line module with a shared tunnel-server port, the available bandwidth for tunnel services is limited to 0.5 Gbps per module. When you configure the ES2 4G line module with a shared tunnel-server port, the available bandwidth for tunnel services is limited to 0.8 Gbps per module.

For information about configuring tunnel services on dedicated and shared tunnel-server ports, see the *Managing Tunnel-Service and IPsec-Service Interfaces* chapter in the *JunosE Physical Layer Configuration Guide*.

For information about line modules supported by the LAC and LNS and the type of support each module type receives, see *ERX Module Guide, Appendix A, Module Protocol Support*.

## E120 Router and E320 Router

To use an LNS on an E120 router or an E320 router, you must install an ES2 4G line module (LM) or an ES2 10G ADV LM with an ES2-S1 Service I/O adapter (IOA). With the ES2 4G LM, it is also possible to use an LNS with an IOA that supports the use of shared tunnel-server ports. For information about installing modules in these routers, see the *E120 and E320 Hardware Guide*.

The combination of an ES2 4G LM or an ES2 10G ADV LM with an ES2-S1 Service IOA provides a dedicated tunnel-server port that is always configured on the IOA. Unlike SMs, the ES2 4G LM and the ES2 require the ES2-S1 Service IOA to condition it to receive and transmit data to other line modules. The ES2-S1 Service IOA also does not have ingress or egress ports. The ES2 10G ADV LM with the ES2-S1 Service IOA supports L2TP LNS functionality, which supports IPv4 as well as IPv6 encapsulated within PPP and L2TP over IPv4.

You can also create tunnels on IOAs that support shared tunnel-server ports. You can configure (provision) a shared tunnel-server port to use a portion of the bandwidth of the IOA to provide tunnel services. For a list of the IOAs that support shared tunnel-server ports, see the *E120 and E320 Module Guide*.

For information about IOAs that are supported by the LAC and LNS and the type of support each module type receives, see *E120 and E320 Module Guide, Appendix A, IOA Protocol Support*.

### Related Documentation

- [L2TP Overview on page 321](#)
- [Implementing L2TP on page 323](#)
- [L2TP Platform Considerations on page 326](#)
- [L2TP References on page 329](#)

## Sessions and Tunnels Supported

---

The E120 and E320 routers support 60,000 L2TP sessions, the ERX1440 router supports 32,000 L2TP sessions, and all other E Series routers support a maximum of 16,000 L2TP sessions. The following guidelines apply:

- On all E Series routers

The SM and the ES2-S1 Service IOA both support the termination of 16,000 LNS sessions per module. Therefore, if you want to apply input or output policies and QoS shapers to all of the available LNS sessions, you can only terminate a maximum of 8000 sessions per module.

- On the E120 router, E320 router, and the ERX1440 router

You can create a systemwide maximum of 60,000 sessions per E120 or E320 router or 32,000 sessions per ERX1440 router. The maximum number of tunnels supported on the E120, E320, and ERX1440 routers are 16,000, 32,000, and 8,000 respectively. The maximum session limit is spread in any combination across tunnels in the router. For a router that operates as an LAC for some tunnels and as an LNS for others, keep in mind the following conditions:

- Combined total of LAC and LNS tunnels should not exceed the maximum tunnel limit.
  - Combined total of LAC and LNS sessions should not exceed the maximum session limit.
- On all E Series routers except the ERX1440 router, E120 router, and the E320 router
- You can create a systemwide maximum of 16,000 sessions spread in any combination across a maximum of 8000 tunnels shared between an LAC and an LNS. For a router that operates as an LAC for some tunnels and as an LNS for others, keep in mind the following conditions:
- Combined total of LAC and LNS tunnels should not exceed 8,000 tunnels.
  - Combined total of LAC and LNS sessions should not exceed 16,000 sessions.



**NOTE:** In previous releases, the JunosE Software required that you use the `license l2tp-session` command to configure a license to enable support for the maximum allowable L2TP sessions on ERX1440 routers, E120 routers, and E320 routers. The `license l2tp-session` command still appears in the CLI, but it has no effect on the actual enforced limit. The reported license limit is 60,000. The `show license l2tp-session` command also still appears in the CLI.

---

- To obtain the maximum number of ingress and egress policy attachments supported for L2TP sessions, see *JunosE Release Notes, Appendix A, System Maximums*.



- Related Documentation**
- [L2TP Overview on page 321](#)
  - [Implementing L2TP on page 323](#)
  - [L2TP Module Requirements on page 326](#)

## L2TP References

---

For more information about L2TP, see the following resources:

- RFC 2661—Layer Two Tunneling Protocol “L2TP” (August 1999)
- RFC 3145—L2TP Disconnect Cause Information (July 2001)
- Fail Over extensions for L2TP “failover”—draft-ietf-l2tpext-failover-06.txt (April 2006 expiration)
- RFC 4951—Fail Over Extensions for Layer 2 Tunneling Protocol (L2TP) “failover” (August 2007)

For information about L2TP high availability support, see the *Managing High Availability* chapter in the *JunosE System Basics Configuration Guide*.

For information about setting up policy-based routing features for L2TP, such as rate limit profiles, classifier control lists, and policy lists, see the *JunosE Policy Management Configuration Guide*.

For information about creating and attaching QoS profiles to L2TP sessions, see the *JunosE Quality of Service Configuration Guide*.

For information about how to secure Layer 2 Tunneling Protocol (L2TP) tunnels with IP Security (IPsec) on your E Series router, see the *Securing L2TP and IP Tunnels with IPsec* chapter in the *JunosE IP Services Configuration Guide*.

- Related Documentation**
- [L2TP Overview on page 321](#)
  - [L2TP Platform Considerations on page 326](#)



## CHAPTER 13

# Configuring an L2TP LAC

An L2TP access concentrator (LAC) receives packets from a remote client and forwards them to an L2TP network server (LNS) on a remote network. You can configure your E Series router to function as an LAC.

This chapter includes the following topics that provide information for configuring an L2TP LAC on the E Series router:

- [LAC Configuration Prerequisites on page 332](#)
- [Modifying L2TP LAC Default Settings for Managing Destinations, Tunnels, and Sessions on page 332](#)
- [Generating UDP Checksums in Packets to L2TP Peers on page 333](#)
- [Specifying a Destruct Timeout for L2TP Tunnels and Sessions on page 334](#)
- [Preventing Creation of New Destinations, Tunnels, and Sessions on page 334](#)
- [Shutting Down Destinations, Tunnels, and Sessions on page 336](#)
- [Specifying the Number of Retransmission Attempts on page 337](#)
- [Controlling the Transmission of L2TP Hello Messages to L2TP Peers in Control Channels on page 337](#)
- [Configuring Calling Number AVP Formats on page 338](#)
- [Mapping a User Domain Name to an L2TP Tunnel Overview on page 347](#)
- [Mapping User Domain Names to L2TP Tunnels from Domain Map Tunnel Mode on page 348](#)
- [Mapping User Domain Names to L2TP Tunnels from Tunnel Group Tunnel Mode on page 352](#)
- [Transmission of the Subscriber Access Interface Speed to LNS Using the RX Connect-Speed AVP on page 354](#)
- [Configuring the RX Speed on the LAC on page 355](#)
- [Managing the L2TP Destination Lockout Process on page 357](#)
- [Managing Address Changes Received from Remote Endpoints on page 360](#)
- [Configuring LAC Tunnel Selection Parameters on page 361](#)

## LAC Configuration Prerequisites

Before you begin configuring the router as an LAC, perform the following steps:

1. Create a virtual router.

```
host1(config)#virtual-router west
```

2. Assign a router ID IP address, such as that for a loopback interface, to the virtual router. This address must be reachable by the L2TP peer.

```
host1:west(config)#ip router-id 10.10.45.3
```



**CAUTION:** You must explicitly assign a router ID to a virtual router rather than using a dynamically assigned router ID. A fixed ID is required because every time the ID changes, L2TP must disconnect all existing tunnels and sessions that use the old ID. If you use a dynamically assigned router ID, the value can change without warning, leading to failure of all L2TP tunnels and sessions. Also, the router could dynamically assign a router ID that is not reachable by the L2TP peer, causing a complete failure of L2TP. You must set the router ID even if you specified a source address in the domain map or a local address in the host profile.

3. When configuring the router as a LAC, configure the router or virtual router for Broadband Remote Access Server (B-RAS).



**NOTE:** If you are using shared tunnel-server ports, you must configure the shared tunnel-server ports before you configure Layer 2 Tunneling Protocol (L2TP) network server (LNS) support. You use the `tunnel-server` command in Global Configuration mode to specify the physical location of the shared tunnel-server port that you want to configure.

See *JunosE Physical Layer Configuration Guide* for additional information about the `tunnel-server` command and shared tunnel-server ports.

- Related Documentation**
- *virtual-router*
  - *ip router-id*

## Modifying L2TP LAC Default Settings for Managing Destinations, Tunnels, and Sessions

Configuring an E Series router for B-RAS enables the router to operate as an LAC with default settings. You can modify the default settings as follows:

- Enable the verification of data integrity via UDP.
- Specify the time period for which the router maintains dynamic destinations, tunnels, or sessions after termination.



**NOTE:** The previous two operations also apply to an LNS, however there is no default configuration that enables the LNS.

When the router is established as an LAC or LNS and is creating destinations, tunnels, and sessions, you can manage them as follows:

- Prevent the creation of new sessions, tunnels, and destinations.
- Close and reopen all or selected destinations, tunnels, and sessions.
- Configure drain timeout operations, which control the amount of time a disconnected LAC tunnel waits before restarting after receiving a restart request.
- Configure how many times the router retries a transmission if the initial attempt is unsuccessful.



**NOTE:** All the commands in this section apply to both the LAC and the LNS.

#### Related Documentation

- [Generating UDP Checksums in Packets to L2TP Peers on page 333](#)
- [Specifying a Destruct Timeout for L2TP Tunnels and Sessions on page 334](#)
- [Preventing Creation of New Destinations, Tunnels, and Sessions on page 334](#)
- [Shutting Down Destinations, Tunnels, and Sessions on page 336](#)
- [Specifying the Number of Retransmission Attempts on page 337](#)

## Generating UDP Checksums in Packets to L2TP Peers

You can configure the router to generate a UDP data integrity checksum in data packets sent to an L2TP peer. The router always uses UDP checksums during transmission and reception of L2TP control packets. Generation of checksums is disabled by default.

- To enable generation of UDP checksums:

```
host1(config)#l2tp checksum
```



**NOTE:** This command does not affect the way the router checks the UDP data integrity checksum in L2TP data packets that are received from an L2TP peer. The router checks all non-zero received checksums and discards the packet if a data integrity problem is detected.

L2TP checksum generation support is available on an ES2 10G Uplink LM and an ES2 4G LM only. It is not supported on an ES2 10G LM and an ES2 10G ADV LM. If an ES2 10G LM or an ES2 10G ADV LM is present when L2TP checksum is enabled, the checksum is not calculated and its value is set to zero.

**Related Documentation** • [l2tp checksum](#)

---

## Specifying a Destruct Timeout for L2TP Tunnels and Sessions

---

You can specify the maximum time period, in the range 10–3600 seconds (1 hour), for which the router attempts to maintain dynamic destinations, tunnels, and sessions after they have been destroyed. The router uses a timeout of 600 seconds by default.

This command facilitates debugging and other analysis by saving underlying memory structures after the destination, tunnel, or session is terminated.

Any specific dynamic destination, tunnel, or session may not be maintained for this entire time period if the resources must be reclaimed early to allow new tunnels to be established.

When a subscriber is terminated, the server port that hosted the subscriber session is released after the dynamic interface destruct timeout is exceeded. The server port that is released is available for a new incoming-call request (ICRQ) packet that the LAC sends to the LNS. Until the time any server port is available to be used for a new incoming call, new ICRQ packets are denied because of a lack of system resources.



**TIP:** If you use the `l2tp destination lockout timeout` command to configure an optional lockout timeout, always configure the destruct timeout to be longer than the lockout timeout. The destruct timeout overrides the lockout timeout—when the destruct timeout expires, all information about the locked out destination is deleted, including the lockout timeout and lockout test settings. See [“Managing the L2TP Destination Lockout Process” on page 357](#).

- 
- To specify a destruct timeout:

```
host1(config)#l2tp destruct-timeout 1200
```

**Related Documentation** • [l2tp destruct-timeout](#)

---

## Preventing Creation of New Destinations, Tunnels, and Sessions

---

You can configure several L2TP drain operations, which determine how the router creates new L2TP destinations, tunnels, and sessions. You can manage the following features:

1. [Preventing Creation of New Destinations, Tunnels, and Sessions on the Router on page 335](#)
2. [Preventing Creation of New Tunnels and Sessions at a Destination on page 335](#)
3. [Preventing Creation of New Sessions for a Tunnel on page 335](#)
4. [Specifying a Drain Timeout for a Disconnected Tunnel on page 335](#)

## Preventing Creation of New Destinations, Tunnels, and Sessions on the Router

You use the **l2tp drain** command to prevent the creation of new destinations, tunnels, and sessions on the router.

The **l2tp drain** command and the **l2tp shutdown** command both affect the administrative state of L2TP on the router. Although each command has a different effect, the **no** version of each command is equivalent. Each command's **no** version leaves L2TP in the enabled state.

- To prevent the creation of new destinations, tunnels, and sessions:

```
host1(config)#l2tp drain
```

## Preventing Creation of New Tunnels and Sessions at a Destination

You use the **l2tp drain destination** command to prevent the creation of new tunnels and sessions at a specific destination.

The **l2tp drain destination** command and the **l2tp shutdown destination** command both affect the administrative state of L2TP for the destination. Although each command has a different effect, the **no** version of each command is equivalent. Each command's **no** version leaves L2TP in the enabled state.

- To prevent the creation of new tunnels and sessions at the specified destination:

```
host1(config)#l2tp drain destination ip 172.31.1.98
```

## Preventing Creation of New Sessions for a Tunnel

Use the **l2tp drain tunnel** command to prevent the creation of new sessions for a tunnel.

The **l2tp drain tunnel** command and the **l2tp shutdown tunnel** command both affect the administrative state of L2TP for the tunnel. Although each command has a different effect, the **no** version of each command is equivalent. Each command's **no** version leaves L2TP in the enabled state.

- To prevent the creation of new sessions for a specific tunnel:

```
host1(config)#l2tp drain tunnel virtual-router default ip 172.31.1.98 isp.com
```

## Specifying a Drain Timeout for a Disconnected Tunnel

Use the **l2tp tunnel short-drain-timeout** command to specify the amount of time a disconnected LAC L2TP tunnel waits before restarting after it receives a restart request.

You can specify a drain timeout in the range 0–31 seconds. This feature enables the router to restart tunnels more quickly than the standard 31-second drain time specified by RFC-2661. By default, the router uses a short-drain timeout of 2 seconds.

- To specify the short-drain timeout:

```
host1(config)#l2tp tunnel short-drain-timeout 12
```

## Shutting Down Destinations, Tunnels, and Sessions

---

You can configure how the router shuts down L2TP destinations, tunnels, and sessions. You can specify the following shutdown methods, which also prevent the creation of new tunnels:

1. [Closing and Preventing Existing and New Destinations, Tunnels, and Sessions on the Router on page 336](#)
2. [Closing and Preventing Existing and New Tunnels and Sessions for a Destination on page 336](#)
3. [Closing and Preventing Existing and New Sessions in a Specific Tunnel on page 336](#)
4. [Closing a Specific Session on page 337](#)

### Closing and Preventing Existing and New Destinations, Tunnels, and Sessions on the Router

You use the **l2tp shutdown** command to close all existing destinations, tunnels, and sessions, and to prevent the creation of new destinations, tunnels, and sessions on the router.

If an SCCRQ message is received when the **l2tp shutdown** command is configured, then a StopCCN packet along with the appropriate Result Code AVP (Result Code value set to 2 and Error Code to 4) is sent to the LAC to indicate that the LNS is shut down.

The **l2tp shutdown** and **l2tp drain** commands affect the administrative state of L2TP on the router. Although each command has a different effect, the **no** version of each command leaves L2TP in the enabled state.

- To close all destinations, tunnels, and sessions on the router:

```
host1(config)#l2tp shutdown
```

### Closing and Preventing Existing and New Tunnels and Sessions for a Destination

You use the **l2tp shutdown destination** command to close all existing tunnels and sessions for a destination and to prevent the creation of tunnels and sessions for that destination.

The **l2tp shutdown destination** and **l2tp drain destination** commands affect the administrative state of L2TP for the destination. Although each command has a different effect, the **no** version of each command leaves L2TP in the enabled state.

- To close tunnels and sessions, and prevent the creation of new tunnels and sessions for the specified destination:

```
host1(config)#l2tp shutdown destination 1
```

### Closing and Preventing Existing and New Sessions in a Specific Tunnel

You use the **l2tp shutdown tunnel** command to close all sessions in a tunnel and to prevent the creation of sessions in a tunnel.



The **l2tp shutdown tunnel** command and the **l2tp drain tunnel** command both affect the administrative state of L2TP for the tunnel. Although each command has a different effect, the **no** version of each command is equivalent. Each command's **no** version leaves L2TP in the enabled state.

- To close all existing sessions in a specific tunnel and prevent creation of new sessions:

```
host1(config)#l2tp shutdown tunnel 1/isp.com
```

## Closing a Specific Session

You use the **l2tp shutdown session** command to close the specified session.

- To close a specific session:

```
host1(config)#l2tp shutdown session 1/1/1
```

## Specifying the Number of Retransmission Attempts

You can specify the number of retransmission attempts the router uses for tunnels, in the range 2–30. By default, the router uses a retry count of 5.

Use the **established** keyword to apply the retry count only to established tunnels. Use the **not-established** keyword to apply the retry count only to tunnels that are not established. If you do not include a keyword, the router applies the retry count to both established and nonestablished tunnels.

- To configure the number of retransmission attempts:

```
host1(config)#l2tp retransmission 4 established
```

If you perform a stateful SRP switchover on an LNS device, we recommend that you configure the maximum number of retransmission attempts as 10, although the default number of attempts is 5. This recommendation applies for all types of L2TP peer resynchronization methods configured for LNS devices.

### Related Documentation

- *l2tp retransmission*

## Controlling the Transmission of L2TP Hello Messages to L2TP Peers in Control Channels

In JunosE implementation, when a control channel is up, the E Series router that is acting as an L2TP access concentrator (LAC) or L2TP network server (LNS) transmits a Hello packet by default on a lapse of 60 seconds since the transmission or reception of the last control packet in the control channel. The Layer 2 Tunneling Protocol (L2TP) peer of the LAC or LNS is expected to acknowledge the reception of the Hello message with a Zero Length Body (ZLB) message. The LAC or LNS terminates the control channel if the Hello packet is not acknowledged. However, there could be some implementations where the peer is incapable of responding to the Hello messages with ZLBs, which causes the LAC or LNS to unintentionally terminate the control channel. In such cases, you can use the **l2tp disable tunnel-hello** command to disable the transmission of Hello messages from the LAC or LNS.



**NOTE:** When the transmission of Hello messages is disabled, the following effects are expected:

- If a Point-to-Point Protocol (PPP) keepalive message is also disabled and the L2TP peer goes out of reach, the control connection cannot be closed at the LAC or LNS end because no explicit control message is received by the LAC or LNS from its peer. This causes the connection to use the LAC or LNS resources unnecessarily. If the maximum number of connections were in use before the peer went out of reach, the peer might not be able to reconnect after it comes up because all the L2TP resources are used up by the previous connections. The resultant behavior depends on the peer L2TP implementation.
- JunosE Silent Failover (SF)–based control connection recovery is dependent on the transmission of Hello messages to synchronize sequence numbers with the L2TP peer after failover. With the transmission of Hello messages disabled, SF-based recovery cannot be performed.

To avoid the transmission of Hello messages to the L2TP peer:

- Issue the **l2tp disable tunnel-hello** command in Global Configuration mode.

```
host1(config)#l2tp disable tunnel-hello
```

Use the **no** version to restore the transmission of Hello messages.



**NOTE:** Execution of the **no l2tp disable tunnel-hello** command resumes the transmission of Hello messages in the tunnel and results in the termination of the control connection if the L2TP peer did not support the transmission of Hello messages.

**Related  
Documentation**

- *l2tp retransmission*

---

## Configuring Calling Number AVP Formats

---

The E Series LAC generates L2TP Calling Number AVP 22 for incoming-call request (ICRQ) packets that the LAC sends to the LNS. By default, the E Series LAC generates the Calling Number AVP 22 in descriptive format.

You can also prevent the E Series LAC from sending the Calling Number AVP in ICRQ packets.



**NOTE:** You cannot change the L2TP Calling Number AVP on tunnel switched interfaces.

You use the **aaa tunnel calling-number-format** command to configure the router to generate AVP 22 in any of the following formats. Agent-circuit-id is suboption 1 of the tags supplied by the PPPoE intermediate agent from the DSLAM. Agent-remote-id is suboption 2.

- **descriptive**—This is the default format, and includes the following elements:  
`<interface ID> <delimiter> <UID> <delimiter> <interface description> <delimiter> <connect info> <delimiter> <PPPoE description>`
- **descriptive include-agent-circuit-id**—This format includes the following elements:  
`<interface ID> <delimiter> <UID> <delimiter> <interface description> <delimiter> <connect info> <delimiter> <PPPoE description> <delimiter> <agent-circuit-id>`
- **descriptive include-agent-circuit-id include-agent-remote-id**—This format includes the following elements:  
`<interface ID> <delimiter> <UID> <delimiter> <interface description> <delimiter> <connect info> <delimiter> <PPPoE description> <delimiter> <agent-circuit-id> <delimiter> <agent-remote-id>`
- **descriptive include-agent-remote-id**—This format includes the following elements:  
`<interface ID> <delimiter> <UID> <delimiter> <interface description> <delimiter> <connect info> <delimiter> <PPPoE description> <delimiter> <agent-remote-id>`
- **fixed**—This format is similar to the fixed format of RADIUS attribute 31 (Calling-Station-Id). If you set up the router to generate the Calling Number AVP in fixed format, the router formats the AVP to use a fixed format of up to 15 characters consisting of all ASCII fields, as follows (the maximum number of characters for each field is shown in brackets):
  - For ATM interfaces:  
`<system name [4]> <slot [2]> <port [1]> <VPI [3]> <VCI [5]>`
  - For Ethernet interfaces:  
`<system name [4]> <slot [2]> <port [1]> <VLAN [8]>`
  - Format for serial interfaces:  
`<system name [4]> <slot [2]> <port [1]> <O [8]>`
  - **Example**—The following command configures the L2TP Calling Number AVP in fixed format:  

```
host1(config)#aaa tunnel calling-number-format fixed
```

For example, when you configure this L2TP Calling Number AVP format on an E320 Broadband Services Router for an ATM interface on system name eastern, slot 14, adapter 1, port 2, VCI 3, and VPI 4, the virtual router displays the format in ASCII as '14' '2' '003' '00004'. The adapter number does not appear in this format.
- **fixed-adapter-embedded**—If you set up the router to generate the L2TP Calling Number AVP in fixed-adapter-embedded format, the router formats the AVP to use a fixed format of up to 15 characters consisting of all ASCII fields with a 1-byte *slot* field, 1-byte *adapter* field, and 1-byte *port* field:

- Format for ATM interfaces:  
*systemName* (up to 4 bytes) *slot* (1 byte) *adapter* (1 byte)  
*port* (1 byte) *VPI* (3 bytes) *VCI* (5 bytes)
- Format for Ethernet interfaces:  
*systemName* (up to 4 bytes) *slot* (1 byte) *adapter* (1 byte)  
*port* (1 byte) *VLAN* (8 bytes)
- Format for serial interfaces:  
*systemName* (up to 4 bytes) *slot* (1 byte) *adapter* (1 byte)  
*port* (1 byte) 0 (8 bytes)
- For E120 and E320 Broadband Services Routers, *adapter* is the number of the bay in which the I/O adapter (IOA) resides, either 0 (representing the right IOA bay on the E120 router and the upper IOA bay on the E320 router) or 1 (representing the left IOA bay on the E120 router or the lower IOA bay on the E320 router). For ERX7xx models, ERX14xx models, and ERX310 Broadband Services Routers, which do not use IOAs, *adapter* is always shown as 0.
- Slot numbers 0 through 16 are shown as ASCII characters in the 1-byte slot field according to the following translation:

| Slot Number | ASCII Character | Slot Number | ASCII Character |
|-------------|-----------------|-------------|-----------------|
| 0           | 0               | 9           | 9               |
| 1           | 1               | 10          | A               |
| 2           | 2               | 11          | B               |
| 3           | 3               | 12          | C               |
| 4           | 4               | 13          | D               |
| 5           | 5               | 14          | E               |
| 6           | 6               | 15          | F               |
| 7           | 7               | 16          | G               |
| 8           | 8               | –           | –               |

For example, slot 16 is shown as the ASCII character uppercase G.

- Example—The following command configures the L2TP Calling Number AVP in fixed-adapter-embedded format:

```
host1(config)#aaa tunnel calling-number-format fixed-adapter-embedded
```

For example, when you configure this L2TP Calling Number AVP format on an E320 router for an ATM interface on system name eastern, slot 14, adapter 1, port 2, VCI 3, and VPI 4, the virtual router displays the format in ASCII as 'E' '1' '2' '003' '00004'.

- **fixed-adapter-new-field**—If you set up the router to generate the L2TP Calling Number AVP in fixed-adapter-embedded-new-field format, the router formats the AVP to use a fixed format of up to 17 characters consisting of all ASCII fields with a 2-byte *slot* field, 1-byte *adapter* field, and 2-byte *port* field:
  - Format for ATM interfaces:  
*systemName* (up to 4 bytes) *slot* (2 bytes) *adapter* (1 byte)  
*port* (2 bytes) *VPI* (3 bytes) *VCI* (5 bytes)
  - Format for Ethernet interfaces:  
*systemName* (up to 4 bytes) *slot* (2 bytes) *adapter* (1 byte)  
*port* (2 bytes) *VLAN* (8 bytes)
  - Format for serial interfaces:  
*systemName* (up to 4 bytes) *slot* (2 bytes) *adapter* (1 byte)  
*port* (2 bytes) 0 (8 bytes)
- Slot numbers 0 through 16 are shown as integers in the 2-byte *slot* field.
- Example—The following command configures the L2TP Calling Number AVP in fixed-adapter-new-field format:

```
host1(config)#aaa tunnel calling-number-format fixed-adapter-new-field
```

For example, when you configure this L2TP Calling Number AVP format on an E320 router for an ATM interface on system name eastern, slot 14, adapter 1, port 2, VCI 3, and VPI 4, the virtual router displays the format in ASCII as '14' '1' '02' '003' '00004'.

- **include-agent-circuit-id format**—This format includes the following element:  
`<agent-circuit-id>`
- **include-agent-circuit-id include-agent-remote-id format**—This format includes the following elements:  
`<agent-circuit-id> <delimit> <agent-remote-id>`
- **include-agent-remote-id format**—This format includes the following element:  
`<agent-remote-id>`
- **stacked**—This format includes a 4-byte stacked VLAN (S-VLAN) ID in the fixed, fixed-adapter-embedded, and fixed-adapter-new-field Calling Number AVP formats for Ethernet interfaces. The S-VLAN ID is displayed in decimal format in the range 0–4095. By default, these formats do not include the S-VLAN ID unless you specify the optional **stacked** keyword.



**NOTE:** The use of the **stacked** keyword is not supported for VLAN subinterfaces based on agent-circuit-identifier information, otherwise known as ACI VLANs. When you issue the **aaa tunnel calling-number-format fixed stacked**, **aaa tunnel calling-number-format fixed-adapter-embedded stacked**, or **aaa tunnel calling-number-format fixed-adapter-new-field stacked** command for an ACI VLAN, the values that appear in the 4-byte S-VLAN ID and 4-byte VLAN ID fields are incorrect.

- Format for Ethernet interfaces that use **fixed**:  
*systemName* (up to 4 bytes) *slot* (2 bytes) *port* (1 byte) *S-VLAN* (4 bytes) *VLAN* (4 bytes)
- Format for Ethernet interfaces that use **fixed-adapter-embedded**:  
*systemName* (up to 4 bytes) *slot* (1 byte) *adapter* (1 byte) *port* (1 byte) *S-VLAN* (4 bytes) *VLAN* (4 bytes)
- Format for Ethernet interfaces that use **fixed-adapter-new-field**:  
*systemName* (up to 4 bytes) *slot* (2 bytes) *adapter* (1 byte) *port* (2 bytes) *S-VLAN* (4 bytes) *VLAN* (4 bytes)
- The S-VLAN ID field in the Calling Number AVP is set to 0 (zero) if you do not specify the optional **stacked** keyword, or if you specify the optional **stacked** keyword but the Ethernet interface does not have an S-VLAN ID.
- Example—The following command configures the L2TP Calling Number AVP in fixed-adapter-new-field format for an Ethernet interface with an S-VLAN ID:

```
host1(config)#aaa tunnel calling-number-format fixed-adapter-new-field stacked
```

For example, when you configure this Calling-Station-Id format on an E320 router for an Ethernet interface on system name western, slot 4, adapter 1, port 3, S-VLAN ID 8, and VLAN ID 12, the virtual router displays the format in ASCII as 'west' '04' '1' '03' '0008' '0012'.

Tasks for configuring the L2TP Calling Number AVP 22 include:

- [Calling Number AVP 22 Configuration Tasks on page 342](#)
- [Configuring the Fallback Format on page 343](#)
- [Disabling the Calling Number AVP on page 346](#)

## Calling Number AVP 22 Configuration Tasks

To set up the router to generate Calling Number AVP 22 for an Ethernet interface in fixed format that includes both an S-VLAN ID and a VLAN ID:

1. Set the calling number format of the tunnel to **fixed**, and specify the optional **stacked** keyword to include the S-VLAN ID.

```
host1(config)#aaa tunnel calling-number-format fixed stacked
```

2. Set the format of the RADIUS Calling-Station-Id to **fixed-format**, and specify the optional **stacked** keyword to include the S-VLAN ID.

```
host1(config)#radius calling-station-format fixed-format stacked
```

If you use a RADIUS server to authenticate the L2TP tunnel parameters, you must configure the format for both the L2TP Calling Number AVP 22 (by using the **aaa tunnel calling-number-format** command) and the RADIUS Calling-Station-ID [31] attribute (by using the **radius calling-station-format** command).

However, if you use an AAA domain map to authenticate the L2TP tunnel parameters, you need configure only the L2TP Calling Number AVP 22 format by using the **aaa tunnel**

**calling-number-format** command. You need not configure the format of the RADIUS Calling-Station-ID [31] attribute in this case.

## Configuring the Fallback Format

You can configure a fallback AVP 22 format. The E Series LAC uses the fallback format to generate the L2TP Calling Number AVP 22 in the event that the PPPoE agent ID is null or unavailable. The LAC uses the fallback format only when the configured calling number format includes either or both of the agent-circuit-id and agent-remote-id suboptions.

The calling number format determines what element triggers use of the fallback format, as shown in the following table:

| Calling Number Format                                        | Fallback Trigger                                     |
|--------------------------------------------------------------|------------------------------------------------------|
| agent-circuit-id                                             | agent-circuit-id is empty                            |
| agent-circuit-id include-agent-remote-id                     | Both agent-circuit-id and agent-remote-id are empty. |
| agent-remote-id                                              | agent-remote-id is empty                             |
| descriptive include-agent-circuit-id                         | agent-circuit-id is empty                            |
| descriptive include-agent-circuit-id include-agent-remote-id | Both agent-circuit-id and agent-remote-id are empty. |
| descriptive include-agent-remote-id                          | agent-remote-id is empty                             |

You use the **aaa tunnel calling-number-format-fallback** command to configure the router to generate any of the following fallback AVP 22 formats:

- **descriptive**—This is the default fallback AVP 22 format, and includes the following elements:  
`<interface ID> <delimit> <UID> <delimit> <interface description> <delimit> <connect info> <delimit> <PPPoE description>`
- **fixed**—This format is similar to the fixed format of RADIUS attribute 31 (Calling-Station-Id). If you set up the router to generate the fallback AVP 22 in fixed format, the router formats the AVP to use a fixed format of up to 15 characters consisting of all ASCII fields, as follows (the maximum number of characters for each field is shown in brackets):
  - Fallback format for ATM interfaces:  
`<system name [4]> <slot [2]> <port [1]> <VPI [3]> <VCI [5]>`
  - Fallback format for Ethernet interfaces:  
`<system name [4]> <slot [2]> <port [1]> <VLAN [8]>`
  - Fallback format for serial interfaces:  
`<system name [4]> <slot [2]> <port [1]> <0 [8]>`
  - **Example**—The following command configures the fallback AVP 22 in fixed format:

**host1(config)#aaa tunnel calling-number-format-fallback fixed**

For example, when you configure this fallback format on an E320 router for an ATM interface on system name eastern, slot 14, adapter 1, port 2, VCI 3, and VPI 4, the virtual router displays the format in ASCII as '14' '2' '003' '00004'. The adapter number does not appear in this format.

- **fixed-adapter-embedded**—If you set up the router to generate the fallback AVP 22 in fixed-adapter-embedded format, the router formats the AVP to use a fixed format of up to 15 characters consisting of all ASCII fields with a 1-byte *slot* field, 1-byte *adapter* field, and 1-byte *port* field:
  - Fallback format for ATM interfaces:  
*systemName* (up to 4 bytes) *slot* (1 byte) *adapter* (1 byte)  
*port* (1 byte) *VPI* (3 bytes) *VCI* (5 bytes)
  - Fallback format for Ethernet interfaces:  
*systemName* (up to 4 bytes) *slot* (1 byte) *adapter* (1 byte)  
*port* (1 byte) *VLAN* (8 bytes)
  - Fallback format for serial interfaces:  
*systemName* (up to 4 bytes) *slot* (1 byte) *adapter* (1 byte)  
*port* (1 byte) 0 (8 bytes)
- For E120 routers and E320 routers, *adapter* is the number of the bay in which the I/O adapter (IOA) resides, either 0 (representing the right IOA bay on the E120 router and the upper IOA bay on the E320 router) or 1 (representing the left IOA bay on the E120 router or the lower IOA bay on the E320 router). For ERX7xx models, ERX14xx models, and ERX310 routers, which do not use IOAs, *adapter* is always shown as 0.
- Slot numbers 0 through 16 are shown as ASCII characters in the 1-byte slot field according to the following translation:

| Slot Number | ASCII Character | Slot Number | ASCII Character |
|-------------|-----------------|-------------|-----------------|
| 0           | 0               | 9           | 9               |
| 1           | 1               | 10          | A               |
| 2           | 2               | 11          | B               |
| 3           | 3               | 12          | C               |
| 4           | 4               | 13          | D               |
| 5           | 5               | 14          | E               |
| 6           | 6               | 15          | F               |
| 7           | 7               | 16          | G               |
| 8           | 8               | —           | —               |



For example, slot 16 is shown as the ASCII character uppercase G.

- Example—The following command configures the fallback AVP 22 in fixed-adapter-embedded format:

```
host1(config)#aaa tunnel calling-number-format-fallback fixed-adapter-embedded
```

For example, when you configure this fallback format on an E320 router for an ATM interface on system name eastern, slot 14, adapter 1, port 2, VCI 3, and VPI 4, the virtual router displays the format in ASCII as 'E' '1' '2' '003' '00004'.

- fixed-adapter-new-field—If you set up the router to generate the fallback AVP 22 in fixed-adapter-embedded-new-field format, the router formats the AVP to use a fixed format of up to 17 characters consisting of all ASCII fields with a 2-byte *slot* field, 1-byte *adapter* field, and 2-byte *port* field:
  - Fallback format for ATM interfaces:  
*systemName* (up to 4 bytes) *slot* (2 bytes) *adapter* (1 byte)  
*port* (2 bytes) *VPI* (3 bytes) *VCI* (5 bytes)
  - Fallback format for Ethernet interfaces:  
*systemName* (up to 4 bytes) *slot* (2 bytes) *adapter* (1 byte)  
*port* (2 bytes) *VLAN* (8 bytes)
  - Fallback format for serial interfaces:  
*systemName* (up to 4 bytes) *slot* (2 bytes) *adapter* (1 byte)  
*port* (2 bytes) 0 (8 bytes)
  - Slot numbers 0 through 16 are shown as integers in the 2-byte *slot* field.
  - Example—The following command configures the fallback AVP 22 in fixed-adapter-new-field format:

```
host1(config)#aaa tunnel calling-number-format-fallback fixed-adapter-new-field
```

For example, when you configure this fallback format on an E320 router for an ATM interface on system name eastern, slot 14, adapter 1, port 2, VCI 3, and VPI 4, the virtual router displays the format in ASCII as '14' '1' '02' '003' '00004'.

- stacked—This format includes a 4-byte stacked VLAN (S-VLAN) ID in the fixed, fixed-adapter-embedded, and fixed-adapter-new-field fallback AVP 22 formats for Ethernet interfaces. The S-VLAN ID is displayed in decimal format in the range 0–4095. By default, these formats do not include the S-VLAN ID unless you specify the optional **stacked** keyword.



**NOTE:** The use of the **stacked** keyword is not supported for VLAN subinterfaces based on agent-circuit-identifier information, otherwise known as ACI VLANs. When you issue the **aaa tunnel calling-number-format-fallback fixed stacked**, **aaa tunnel calling-number-format-fallback fixed-adapter-embedded stacked**, or **aaa tunnel calling-number-format-fallback fixed-adapter-new-field stacked** command for an ACI VLAN, the values that appear in the 4-byte S-VLAN ID and 4-byte VLAN ID fields are incorrect.

- Fallback format for Ethernet interfaces that use **fixed**:  
*systemName* (up to 4 bytes) *slot* (2 bytes) *port* (1 byte) *S-VLAN* (4 bytes) *VLAN* (4 bytes)
- Fallback format for Ethernet interfaces that use **fixed-adapter-embedded**:  
*systemName* (up to 4 bytes) *slot* (1 byte) *adapter* (1 byte) *port* (1 byte) *S-VLAN* (4 bytes) *VLAN* (4 bytes)
- Fallback format for Ethernet interfaces that use **fixed-adapter-new-field**:  
*systemName* (up to 4 bytes) *slot* (2 bytes) *adapter* (1 byte) *port* (2 bytes) *S-VLAN* (4 bytes) *VLAN* (4 bytes)
- The S-VLAN ID field in the fallback AVP 22 is set to 0 (zero) if you do not specify the optional **stacked** keyword, or if you specify the optional **stacked** keyword but the Ethernet interface does not have an S-VLAN ID.
- Example—The following command configures the fallback AVP 22 in fixed-adapter-new-field format for an Ethernet interface with an S-VLAN ID:

```
host1(config)#aaa tunnel calling-number-format-fallback fixed-adapter-new-field stacked
```

For example, when you configure this fallback format on an E320 router for an Ethernet interface on system name western, slot 4, adapter 1, port 3, S-VLAN ID 8, and VLAN ID 12, the virtual router displays the format in ASCII as 'west' '04' '1' '03' '0008' '0012'.

## Disabling the Calling Number AVP

You can use the **l2tp disable calling-number-avp** command to prevent the E Series LAC from sending the Calling Number AVP in ICRQ packets. You use this command in special situations where you do not want the LAC to send this AVP.

- To prevent the LAC from sending the Calling Number AVP:

```
host1(config)#l2tp disable calling-number-avp
```

For more information about setting up the router to generate Calling Number AVP 22 in a format that includes either or both of the agent-circuit-id and agent-remote-id suboptions of the tags supplied by the PPPoE intermediate agent, see *Configuring PPPoE Remote Circuit ID Capture* in the *JunosE Link Layer Configuration Guide*.

### Calling Number AVP 22 Configuration Examples

The following examples show how you can synchronize the contents of RADIUS Calling-Station-Id (Attribute 31) and L2TP Calling-Number (AVP 22).

- To send the PPPoE agent-circuit-id in RADIUS Attribute 31 and L2TP AVP 22 and specify that the fixed format is used when the PPPoE agent-circuit-id is unavailable, issue the following commands:

```
host1(config)#radius calling-station-format fixed-format
host1(config)#radius remote-circuit-id-delimiter #
host1(config)#radius override calling-station-id remote-circuit-id
host1(config)#radius remote-circuit-id-format agent-circuit-id
host1(config)#aaa tunnel calling-number-format include-agent-circuit-id
host1(config)#aaa tunnel calling-number-format-fallback fixed
```

- To send the PPPoE agent-circuit-id and agent-remote-id in RADIUS Attribute 31 and L2TP AVP 22 and specify that the fixed format is used when both PPPoE agent-circuit-id and agent-remote-id are unavailable, issue the following commands:

```
host1(config)#radius calling-station-format fixed-format
host1(config)#radius remote-circuit-id-delimiter #
host1(config)#radius override calling-station-id remote-circuit-id
host1(config)#radius remote-circuit-id-format agent-circuit-id agent-remote-id
host1(config)#aaa tunnel calling-number-format include-agent-circuit-id
include-agent-remote-id
host1(config)#aaa tunnel calling-number-format-fallback fixed
```

## Mapping a User Domain Name to an L2TP Tunnel Overview

The router uses either the local database related to the domain name or a RADIUS server to determine whether to terminate or tunnel PPP connections.

For information about setting up RADIUS to provide this mapping, see the *Configuring Remote Access* chapter.

For a given domain map, you can choose one of two methods to map the domain to an L2TP tunnel locally on the router:

- Configure tunnels for a domain map and then define tunnel attributes from Domain Map Tunnel configuration mode.
- Configure a tunnel group and then define the attributes for its tunnels from Tunnel Group Tunnel Configuration mode. Use this method only when no tunnels are currently defined for the domain map from Domain Map Tunnel configuration mode. By default, tunnel groups are not assigned to the domain map.

After configuring a tunnel group and the attributes for its tunnels, you can assign the tunnel group to the domain map from Domain Map mode. The tunnel group reference in the domain map is used instead of tunnel definitions configured from Domain Map Tunnel configuration mode.

The RADIUS server can reference tunnel groups through the RADIUS Tunnel Group [26-64] attribute. The advantages of RADIUS support for tunnel groups are:

- The RADIUS server can maintain a single tunnel group attribute associated with each user instead of sets of tunnel attributes for each user.

- The RADIUS server can authenticate users before attempting to establish tunnels.

You can configure up to 31 tunnel definitions for an L2TP subscriber using either AAA domain maps or RADIUS returned values. Each tunnel definition contains both fixed-length and variable-length tunnel attributes. All tunnel definitions and their attributes that are stored in AAA are mirrored in a single transaction. When the size of the mirrored storage transaction exceeds 9866 bytes, the router disables stateful SRP switchover (high availability).

The size of the transaction can exceed 9866 bytes when you configure all the variable length tunnel attributes of more than 17 tagged tunnel definitions, using either RADIUS or domain maps, to their maximum values. When the size of a transaction exceeds 9866 bytes, the router now mirrors the tunnel definitions in a different transaction. As a result, stateful SRP switchover is not disabled when you configure all the variable length tunnel attributes of all 31 tunnel definitions to their maximum values or when the RADIUS server sends tunnel attributes whose length exceeds the maximum length.

**Related  
Documentation**

- [Mapping User Domain Names to L2TP Tunnels from Domain Map Tunnel Mode on page 348](#)
- [Mapping User Domain Names to L2TP Tunnels from Tunnel Group Tunnel Mode on page 352](#)

---

## Mapping User Domain Names to L2TP Tunnels from Domain Map Tunnel Mode

---

To map a domain to an L2TP tunnel locally on the router from Domain Map Tunnel mode, perform the following steps:

1. Specify a domain name and enter Domain Map Configuration mode:

```
host1(config)#aaa domain-map westford.com
host1(config-domain-map)#
```

2. Specify a virtual router; in this case, the *default* router is specified.

```
host1(config-domain-map)#router-name default
```

3. Specify a tunnel to configure and enter Domain Map Tunnel Configuration mode:

```
host1(config-domain-map)#tunnel 3
```

4. Specify the LNS endpoint address of a tunnel.

```
host1(config-domain-map-tunnel)#address 192.0.2.13
```

5. (Optional) Assign a tunnel group to the domain map. You can assign a tunnel group only when no tunnels are currently defined for the domain map from AAA Domain Map Tunnel mode.

```
host1(config-domain-map)#tunnel group storm
```

6. Specify a preference for the tunnel.

You can specify up to eight levels of preference, and you can assign the same preference to a maximum of 31 tunnels. When you define multiple preferences for a destination, you increase the probability of a successful connection.

```
host1(config-domain-map-tunnel)#preference 5
```

7. (Optional) Specify an authentication password for the tunnel.

```
host1(config-domain-map-tunnel)#password temporary
```



**NOTE:** If you specify a password for the LAC, the router requires that the peer (the LNS) authenticate itself to the router. In this case, if the peer fails to authenticate itself, the tunnel terminates.

8. (Optional) Specify a hostname for the LAC end of the tunnel.

The LAC sends the hostname to the LNS when communicating to the LNS about the tunnel. The hostname can be up to 64 characters (no spaces).

```
host1(config-domain-map-tunnel)#client-name host4
```



**NOTE:** If the LNS does not accept tunnels from unknown hosts, and if no hostname is specified, the LAC uses the router name as the hostname.

9. (Optional) Specify a server name for the LNS.

This name specifies the hostname expected from the peer (the LNS) when you set up a tunnel. When this name is specified, the peer must identify itself with this name during tunnel startup. Otherwise, the tunnel is terminated. The server name can be up to 64 characters (no spaces).

```
host1(config-domain-map-tunnel)#server-name boston
```

10. (Optional) Specify a source IP address for the LAC tunnel endpoint. All L2TP packets sent to the peer use this source address.

```
host1(config-domain-map-tunnel)#source-address 192.0.3.3
```

By default, the router uses the virtual router's router ID as the source address. You can override this behavior for an L2TP tunnel by specifying a source address. If you do specify a source address, use the address of a stable IP interface (for example, a loopback interface). Make sure that the address is configured in the virtual router for this domain map, and that the address is reachable by the peer.

11. Specify a tunnel identification. (The router groups L2TP sessions with the same tunnel identification into the same tunnel.)

```
host1(config-domain-map-tunnel)#identification acton
```

The router groups L2TP sessions with the same tunnel identification into the same tunnel. This occurs only when both the destination (virtual router, IP address) and the ID are the same.

12. Specify the L2TP tunnel type (RADIUS attribute 64, Tunnel-Type). Currently, the only supported value is L2TP.

```
host1(config-domain-map-tunnel)#type l2tp
```

13. Specify a medium type for the tunnel. (L2TP supports only IP version 4 [IPv4].)

```
host1(config-domain-map-tunnel)#medium ipv4
```

14. (Optional) Specify a default tunnel client name.

```
host1(config-domain-map-tunnel)#exit
host1(config-domain-map)#exit
host1(config)#aaa tunnel client-name boxford
```

If the tunnel client name is not included in the tunnel attributes that are returned from the domain map or authentication server, the router uses the default name.

15. (Optional) Specify a default tunnel password.

```
host1(config)#aaa tunnel password 3&92k%b#q4
host1(config)#exit
```

If the tunnel password is not included in the tunnel attributes that are returned from the domain map or authentication server, the router uses the default password.

16. (Optional) Set the format for the tunnel assignment ID that is passed to PPP/L2TP.

The tunnel assignment ID format can be either only assignmentID or clientAuthId + serverAuthId + assignmentId.

```
host1(config)#aaa tunnel assignment-id-format assignmentID
```

If you do not set a tunnel assignment ID, the software sets it to the default (assignmentID). This parameter is only generated and used by the L2TP LAC device.

17. (Optional) Specify whether or not to use the tunnel peer's Nas-Port [5] and Nas-Port-Type [61] attributes.

When enabled, the attribute is supplied by the tunnel peer. When disabled, the attribute is not supplied. Use the **no** version of the command to restore the default, enable.

```
host1(config)#aaa tunnel ignore nas-port enable
host1(config)#aaa tunnel ignore nas-port-type disable
```

18. (Optional) Set up the router to ignore sequence numbers in data packets received on L2TP tunnels.

```
host1(config)#l2tp ignore-receive-data-sequencing
```

This command does not affect the insertion of sequence numbers in packets *sent* from the router.



**BEST PRACTICE:** We recommend that you set up the router to ignore sequence numbers in received data packets if you are using IP reassembly. Because IP reassembly might reorder L2TP packets, out-of-order packets might be dropped when sequence numbers are being used on L2TP data packets.

19. (Optional) Disable the generation of authentication challenges by the local tunnel, so that the tunnel does not send a challenge during negotiation. However, the tunnel does accept and respond to challenges it receives from the peer.

```
host1(config)#l2tp disable challenge
```

20. Verify the L2TP tunnel configuration.

```
host1(config)# show aaa domain-map
```

```
Domain: westford.com; router-name: default; ipv6-router-name: default
```

| Tunnel<br>Tag | Tunnel<br>Peer           | Tunnel<br>Source     | Tunnel<br>Type            | Tunnel<br>Medium | Tunnel<br>Password          | Tunnel<br>Id | Tunnel<br>Client<br>Name |
|---------------|--------------------------|----------------------|---------------------------|------------------|-----------------------------|--------------|--------------------------|
| 3             | 192.168.2.13             | 192.168.3.3          | l2tp                      | ipv4             | temporary                   | acton        | host4                    |
| Tunnel<br>Tag | Tunnel<br>Server<br>Name | Tunnel<br>Preference | Tunnel<br>Max<br>Sessions | Tunnel<br>RWS    | Tunnel<br>Virtual<br>Router |              |                          |
| 3             | boston                   | 5                    | 0                         | system chooses   | vr2                         |              |                          |

```
host1#show aaa tunnel-parameters
```

```
Tunnel password is 3&92k%b#q4
```

```
Tunnel client-name is <NULL>
```

```
Tunnel nas-port-method is none
```

```
Tunnel nas-port ignore disabled
```

```
Tunnel nas-port-type ignore disabled
```

```
Tunnel assignmentId format is assignmentId
```

```
Tunnel calling number format is descriptive
```

#### Related Documentation

- [Mapping User Domain Names to L2TP Tunnels from Tunnel Group Tunnel Mode on page 352](#)
- *aaa domain-map*
- *aaa tunnel assignment-id-format*
- *aaa tunnel client-name*
- *aaa tunnel ignore*
- *aaa tunnel password*
- *address*
- *client-name*
- *identification*
- *l2tp disable challenge*
- *l2tp ignore-receive-data-sequencing*
- *medium ipv4*
- *password*
- *preference*
- *router-name*
- *server-name*
- *source-address*
- *tunnel*
- *tunnel group*

- *type*

## Mapping User Domain Names to L2TP Tunnels from Tunnel Group Tunnel Mode

To map a domain to an L2TP tunnel locally on the router from Tunnel Group Tunnel Configuration mode, perform the following steps:

1. Specify an AAA tunnel group and change the mode to Tunnel Group Tunnel Configuration mode. From Tunnel Group Tunnel Configuration mode, you can add up to 31 tunnel definitions.

```
host1(config)#aaa tunnel-group westford
host1(config-tunnel-group)#
```

2. Specify a tunnel to configure and enter Tunnel Group Tunnel Configuration mode:

```
host1(config-tunnel-group)#tunnel 3
host1(config-tunnel-group-tunnel)#
```

3. Specify a virtual router; in this case, the *default* router is specified.

```
host1(config-tunnel-group-tunnel)#router-name default
```

4. Specify the LNS endpoint address of a tunnel.

```
host1(config-tunnel-group-tunnel)#address 192.0.2.13
```

5. Specify a preference for the tunnel.

You can specify up to eight levels of preference, and you can assign the same preference to a maximum of 31 tunnels. When you define multiple preferences for a destination, you increase the probability of a successful connection.

```
host1(config-tunnel-group-tunnel)#preference 5
```

6. (Optional) Specify an authentication password for the tunnel.

```
host1(config-tunnel-group-tunnel)#password temporary
```



**NOTE:** If you specify a password for the LAC, the router requires that the peer (the LNS) authenticate itself to the router. In this case, if the peer fails to authenticate itself, the tunnel terminates.

7. (Optional) Specify a hostname for the LAC end of the tunnel.

The LAC sends the hostname to the LNS when communicating to the LNS about the tunnel. The hostname can be up to 64 characters (no spaces).

```
host1(config-tunnel-group-tunnel)#client-name host4.
```



**NOTE:** If the LNS does not accept tunnels from unknown hosts, and if no hostname is specified, the LAC uses the router name as the hostname.

8. (Optional) Specify a server name for the LNS.



This name specifies the hostname expected from the peer (the LNS) when you set up a tunnel. When this name is specified, the peer must identify itself with this name during tunnel startup. Otherwise, the tunnel is terminated. The server name can be up to 64 characters (no spaces).

```
host1(config-tunnel-group-tunnel)#server-name boston
```

9. (Optional) Specify a source IP address for the LAC tunnel endpoint. All L2TP packets sent to the peer use this source address.

By default, the router uses the virtual router's router ID as the source address. You can override this behavior for an L2TP tunnel by specifying a source address. If you do specify a source address, use the address of a stable IP interface (for example, a loopback interface). Make sure that the address is configured in the virtual router for this domain map, and that the address is reachable by the peer.

```
host1(config-tunnel-group-tunnel)#source-address 192.0.3.3
```

10. Specify a tunnel identification.

```
host1(config-tunnel-group-tunnel)#identification acton
```

The router groups L2TP sessions with the same tunnel identification into the same tunnel. This occurs only when both the destination (virtual router, IP address) and the ID are the same.

11. Specify a medium type for the tunnel. (L2TP supports only IP version 4 [IPv4].)

```
host1(config-tunnel-group-tunnel)#medium ipv4
```

12. Specify the L2TP tunnel type (RADIUS attribute 64, Tunnel-Type). Currently, the only supported value is L2TP.

```
host1(config-tunnel-group-tunnel)#type l2tp
```

13. Verify the L2TP tunnel configuration.

```
host1(config)# show aaa domain-map
```

```
Domain: westford.com; router-name: default; ipv6-router-name: default
```

| Tunnel<br>Tag | Tunnel<br>Peer | Tunnel<br>Source | Tunnel<br>Type | Tunnel<br>Medium | Tunnel<br>Password | Tunnel<br>Id | Tunnel<br>Client<br>Name |
|---------------|----------------|------------------|----------------|------------------|--------------------|--------------|--------------------------|
| 3             | 192.168.2.13   | 192.168.3.3      | l2tp           | ipv4             | temporary          | acton        | host4                    |

| Tunnel<br>Tag | Tunnel<br>Server<br>Name | Tunnel<br>Preference | Tunnel<br>Max<br>Sessions | Tunnel<br>RWS  | Tunnel<br>Virtual<br>Router |
|---------------|--------------------------|----------------------|---------------------------|----------------|-----------------------------|
| 3             | boston                   | 5                    | 0                         | system chooses | vr2                         |

```
host1#show aaa tunnel-parameters
```

```
Tunnel password is 3&92k%b#q4
```

```
Tunnel client-name is <NULL>
```

```
Tunnel nas-port-method is none
```

```
Tunnel nas-port ignore disabled
```

```
Tunnel nas-port-type ignore disabled
```

```
tunnel assignmentId format is assignmentId
```

```
aaa tunnel calling number format is descriptive
```

**Related Documentation**

- [Mapping User Domain Names to L2TP Tunnels from Domain Map Tunnel Mode on page 348](#)
- *aaa tunnel-group*
- *address*
- *client-name*
- *identification*
- *medium ipv4*
- *password*
- *preference*
- *router-name*
- *server-name*
- *source-address*
- *tunnel*
- *type*

---

## Transmission of the Subscriber Access Interface Speed to LNS Using the RX Connect-Speed AVP

---

The L2TP access concentrator (LAC) conveys the speed of the subscriber access interface to the L2TP network server (LNS) through the transmit (TX) Connect-Speed attribute-value pair (AVP) [24] and receive (RX) Connect-Speed AVP [38] in the Incoming-Call-Connected message, during the establishment of an L2TP session. By default, the receive speed of the access interface is set equal to the calculated transmit speed and the generation of RX Connect-Speed AVP is suppressed.

The RX Connect-Speed AVP is generated when the receive speed of the access interface is set equal to the calculated transmit speed by issuing the **l2tp rx-connect-speed-when-equal** command. In this scenario, the LAC transmits the same value for transmit and receive connect speeds that are sent to the LNS through the TX Connect-Speed AVP and RX Connect-Speed AVP in the Incoming-Call-Connected message.

The advisory receive speeds of Ethernet VLANs and ATM1483 circuits can be configured using the **vlan advisory-rx-speed** and **atm atm1483 advisory-rx-speed** commands. When the advisory receive rate is configured, the LAC transmits the configured advisory receive rate to LNS through the RX Connect-Speed AVP. If the advisory receive rate is configured, the transmission of the advisory receive rate in the RX Connect-Speed AVP takes precedence over the transmission of the calculated transmit speed in the RX Connect-Speed AVP even when the **l2tp rx-connect-speed-when-equal** command is issued.

The transmission of the layer 2 control (L2C) RAM actual upstream rate to the LNS can be enabled using the **l2tp rx-connect-speed-upstream-rate** command. Since there is

no integration between L2C and L2TP, the actual upstream rate is sent from the L2C to AAA, which sends the rate with tunnel parameters to the LAC during the creation of the L2TP tunnel. The LAC transmits the actual upstream rate to the LNS through the RX Connect-Speed AVP. If the actual upstream rate transmission is enabled, the transmission of the actual upstream rate in the RX Connect-Speed AVP takes precedence over the transmission of the configured advisory receive rate and transmission of the calculated transmit speed in the RX Connect-Speed AVP.

The transmission of the actual upstream rate is applicable only for the interface in which L2C is configured. If the interface does not have L2C configured, the actual upstream rate is not transmitted. Instead, one of the following actions is taken:

- The advisory receive rate is transmitted in the RX Connect-Speed AVP, if the advisory rate is configured.
- The calculated transmit speed is transmitted in the RX Connect-Speed AVP, if the **l2tp rx-connect-speed-when-equal** command is issued and the advisory rate is not configured.
- The RX Connect-Speed AVP value is transmitted as 0, if the advisory rate is not configured and the **l2tp rx-connect-speed-when-equal** command is not issued or the **no l2tp rx-connect-speed-when-equal** command is issued.



**NOTE:** The actual upstream rate is independent of any QoS parameterization since it does not influence any QoS factor.

#### Related Documentation

- [Configuring the RX Speed on the LAC on page 355](#)
- *atm atm1483 advisory-rx-speed*
- *l2tp rx-connect-speed-upstream-rate*
- *l2tp rx-connect-speed-when-equal*
- *vlan advisory-rx-speed*

## Configuring the RX Speed on the LAC

You can configure the E Series LAC to generate the L2TP RX Connect-Speed AVP [38], which is transmitted to the LNS in the Incoming-Call-Connected message. The AVP carries one of the following subscriber access interface speeds based on the configuration:

- L2C RAM actual upstream rate
- Configured advisory receive speed
- Calculated transmit speed

By default, the receive speed is set equal to the calculated transmit speed and the generation of the RX Connect-Speed AVP is suppressed. The AVP can be used to generate the RADIUS Connect-Info attribute [77] on the LNS.

To set up the router to generate the RX Connect-Speed AVP [38], perform all or any one of the following steps:

- Configure the advisory receive speed:



**NOTE:** The configured advisory receive speed is sent in the RX Connect-Speed AVP, only if the generation of the AVP for transmitting the actual upstream rate is disabled.

- On the ATM subinterface:

```
host1(config-subif)#atm atm1483 advisory-rx-speed 2000
```

For more information about configuring the advisory speed, see *Configuring ATM* in the *JunosE Link Layer Configuration Guide*.

- On the VLAN subinterface:

```
host1(config-subif)#vlan advisory-rx-speed 2000
```

- Enable generation of the RX Connect-Speed AVP when the receive speed is set equal to the calculated transmit speed.



**NOTE:** The calculated transmit speed is sent in the RX Connect-Speed AVP only if the advisory receive speed is not configured and the generation of the AVP for transmitting the actual upstream rate is disabled.

```
host1(config)#l2tp rx-connect-speed-when-equal
```

- Enable generation of the RX Connect-Speed AVP when you want to send the L2C RAM actual upstream rate in the AVP.



**NOTE:** The actual upstream rate is sent in the AVP even if the advisory receive speed is configured and the generation of the AVP is enabled for sending the calculated transmit speed.

```
host1(config)#l2tp rx-connect-speed-upstream-rate
```

#### Related Documentation

- [Transmission of the Subscriber Access Interface Speed to LNS Using the RX Connect-Speed AVP on page 354](#)
- `atm atm1483 advisory-rx-speed`
- `l2tp rx-connect-speed-upstream-rate`
- `l2tp rx-connect-speed-when-equal`
- `vlan advisory-rx-speed`

## Managing the L2TP Destination Lockout Process

When multiple sets of tunneling parameters are available, L2TP uses a selection algorithm to choose the best tunnel for subscriber traffic. As part of this selection process, the JunosE Software's L2TP implementation includes a lockout feature in which the router locks out, or disregards, destinations that are assumed to be unavailable.

By default, when a destination becomes unavailable, L2TP locks out that destination for a lockout timeout of 300 seconds (5 minutes). After the lockout timeout expires, L2TP assumes that the destination is now available and includes the destination when performing the selection algorithm.

Tasks to manage the L2TP lockout process include:

1. [Modifying the Lockout Procedure on page 357](#)
2. [Verifying That a Locked-Out Destination Is Available on page 358](#)
3. [Configuring a Lockout Timeout on page 358](#)
4. [Unlocking a Destination that is Currently Locked Out on page 359](#)
5. [Starting an Immediate Lockout Test on page 359](#)

### Modifying the Lockout Procedure

You can optionally configure your own lockout procedure by specifying the lockout timeout you want to use or enabling a lockout test, or both. When the lockout timeout expires, the destination is either immediately unlocked (if lockout testing is not enabled) or begins the lockout test to verify that the destination is available.

L2TP performs the lockout test by attempting to establish a tunnel to the unavailable destination. For the test, L2TP must first obtain the parameters for a tunnel to the destination. If no such tunnel currently exists, L2TP must wait until it receives a new session request that has tunnel parameters for the locked out destination. The destination remains locked out while L2TP waits for the tunnel parameters and becomes available only after successful completion of the lockout test. Therefore, if lockout testing is enabled, the destination is actually locked out longer than the lockout timer you specify.

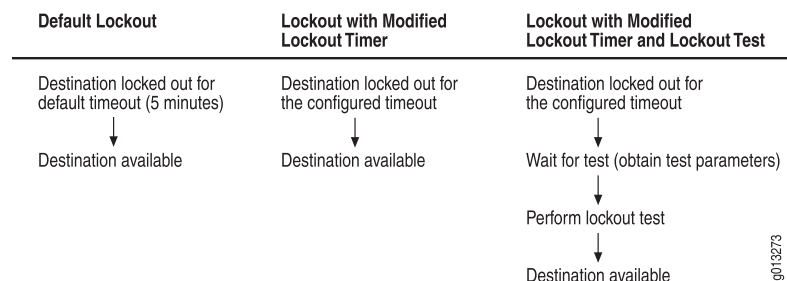


**NOTE:** Always configure the lockout timeout to be shorter than the destruct timeout. The destruct timeout (as described in [“Specifying a Destruct Timeout for L2TP Tunnels and Sessions” on page 334](#)) overrides the lockout timeout—when the destruct timeout expires, all information about the locked out destination is deleted, including the time remaining on the destination's lockout timeout and the requirement to run a lockout test prior to returning the destination to service. As a result, the locked out destination might be returned to service prior to expiration of your configured lockout timeout and without completion of the lockout test you specified.

[Figure 9 on page 358](#) shows how locked-out destinations transition from a locked-out state to available status when using the default lockout configuration, a configuration

that includes a modified lockout timer, and a configuration with both a modified timer and the lockout test.

**Figure 9: Lockout States**



You can use the following commands to manage L2TP destination lockout and configure a lockout process that meets the needs of your network environment:

- Use the **l2tp destination lockout-timeout** command to modify the default lockout timeout period.
- Use the **l2tp destination lockout-test** command to configure L2TP to perform a lockout test, which verifies that a currently locked out destination is now available and to include it in the selection algorithm.
- Use the **l2tp unlock destination** command to force L2TP to immediately unlock the specified locked out destination; the destination is then considered to be available by the selection algorithm. L2TP disregards any time remaining in the existing lockout timeout and also disregards the lockout test (if configured).
- Use the **l2tp unlock-test destination** command to force L2TP to immediately begin the lockout testing procedure for the specified destination; any time remaining in the existing lockout timeout is not taken into account.
- Use the **show l2tp** and **show l2tp destination lockout** commands to view information about the L2TP configuration and statistics.

## Verifying That a Locked-Out Destination Is Available

You can use the **l2tp destination lockout-test** command to configure L2TP to test locked-out destinations; this verifies that a previously locked-out destination is available before the router changes the destination's status.

- To verify the availability of locked out destinations:

```
host1(config)#l2tp destination lockout-test
```

## Configuring a Lockout Timeout

You use the **l2tp destination lockout-timeout** command to configure the amount of time (in seconds) between when an L2TP destination is found to be unavailable and when it is eligible for unlocking. When the timeout period expires, L2TP either begins the lockout test procedure (if configured to do so) or immediately returns the destination to available state.



**BEST PRACTICE:** Always configure the lockout timeout to be shorter than the destruct timeout. The destruct timeout (as described in “[Specifying a Destruct Timeout for L2TP Tunnels and Sessions](#)” on page 334) overrides the lockout timeout—when the destruct timeout expires, all information about the locked out destination is deleted, including the time remaining on the destination's lockout timeout and the requirement to run a lockout test prior to returning the destination to service.

You can specify a lockout timeout in the range 60–3600 seconds (1 minute–1 hour). The router uses a timeout value of 300 seconds by default.

- To configure an L2TP lockout timeout:

```
host1(config)#l2tp destination lockout-timeout 500
```

The new lockout timeout only affects future locked-out destinations; it does not affect destinations that are currently locked out.

## Unlocking a Destination that is Currently Locked Out

You use the **l2tp unlock destination** command to force L2TP to immediately unlock the specified L2TP destination, which is currently locked out and unavailable. L2TP then considers the destination to be available. Any remaining lockout time and the lockout test setting (if configured) are not taken into account.

You must be at privilege level 10 or higher to use this command.

- To unlock a currently locked-out destination:

```
host1(config)#l2tp unlock destination ip 192.168.1.98
```

## Starting an Immediate Lockout Test

You use the **l2tp unlock-test destination** command to force L2TP to immediately start the lockout test for the specified destination—any remaining lockout time for the destination is ignored.

You must be at privilege level 10 or higher to use this command.



**NOTE:** If lockout testing is not configured, this command immediately unlocks the destination and L2TP then considers the destination to be available

- To force an immediate lockout test for a specific destination:

```
host1(config)#l2tp unlock-test destination ip 192.169.110.8
```

## Managing Address Changes Received from Remote Endpoints

A remote endpoint can use the Start-Control-Connection-Reply (SCCRP) packets that it sends to the E Series LAC to change the address that the LAC uses to communicate with the endpoint. By default, the LAC accepts the change and uses the new address to communicate with the endpoint. However, you can configure the LAC to ignore or reject the requested change. Setting up the LAC to ignore address changes in SCCRPs enables the router to construct tunnels with separate receive and transmit addresses and to avoid problems due to a misconfiguration. Three possible configurations are available:

- Default configuration—The E Series LAC accepts the change from the endpoint. The LAC then sends all subsequent packets to, and accepts packets from, the new address.
- Ignore configuration (specified by the **l2tp ignore-transmit-address-change** command)—The LAC continues to send packets to the original address but accepts packets from the new address.

```
host1(config)#l2tp ignore-transmit-address-change
```

Use the **ip-address** or **udp-port** keyword to ignore the specific address component. Omit the keywords to ignore the entire address change in the SCCRPs.

- Reject configuration (specified by the **l2tp reject-transmit-address-change** command)—The LAC sends a Stop-Control-Connection-Notification (StopCCN) to the original address, then terminates the connection to the endpoint.

```
host1(config)#l2tp reject-transmit-address-change ip-address
```

Use the **ip-address** or **udp-port** keyword to reject the specific address component. Omit the keywords to reject the entire address change in the SCCRPs.



**NOTE:** When an L2TP hello message contains a non-zero value in the Reserved Bits field of the L2TP message header, and the LAC rejects the change in the endpoint address by sending a StopCCN to the original address, the Result Code field contains the value of 2 and the Error Code field contains the value of 3. The Result code value denotes a generic error, while the Error code value denotes that one of the field values was out of range or the Reserved Bits field was non-zero in the StopCCN message sent from the LAC to the endpoint.

The reject specification takes precedence over the ignore specification.

The router accepts a change in receive address only once, during the tunnel establishment phase, and only on an SCCRPs packet. Subsequent changes result in the router dropping packets. Any changes do not affect established tunnels.

Use the **show l2tp** command to display the SCCRPs address change configuration.

### Related Documentation

- [l2tp ignore-transmit-address-change](#)
- [l2tp reject-transmit-address-change](#)



## Configuring LAC Tunnel Selection Parameters

This section presents the capabilities of the LAC's tunnel selection process. L2TP allows you to specify:

- Up to 31 destinations for a domain.
- Up to eight levels of preference. Preference indicates the order in which the router attempts to connect to the destinations specified for a domain. Zero (0) is the highest level of preference.
- Up to 31 destinations for a single preference level.

For information about setting up destinations and preference levels for a domain, see [“Mapping a User Domain Name to an L2TP Tunnel Overview” on page 347](#).

When the E Series LAC determines that a PPP session should be tunneled, it selects a tunnel from a set of tunnels associated with either the PPP user or the PPP user's domain. The router provides the following methods for selecting tunnels:

- Tunnel selection failover between preference levels (the default behavior)
  - Tunnel selection failover within a preference level
  - Maximum sessions per tunnel
  - Weighted load balancing
1. [Configuring the Failover Between Preference Levels Method on page 361](#)
  2. [Configuring the Failover Within a Preference Level Method on page 362](#)
  3. [Configuring the Maximum Sessions per Tunnel on page 363](#)
  4. [Configuring the Weighted Load Balancing Method on page 363](#)

### Configuring the Failover Between Preference Levels Method

When a user tries to log into a domain, in the default method, the router attempts to connect to a destination in that domain with the highest preference level. If more than one destination in the preference level is considered reachable, the router randomly selects a destination and attempts to contact it. If the router is unsuccessful, it marks the destination as unreachable and does not try to connect to that destination for five minutes. The router then moves to the next lower preference level and repeats the process. The router makes up to eight attempts to connect to a destination for a domain—one attempt for each preference level.

If all destinations at a preference level are marked as unreachable, the router chooses the destination that failed first and tries to make a connection. The key is to understand that the router chooses a single destination at each level of preference, even if all destinations have recently failed. Thus the 5-minute timer normally used to reinstate failed destinations is ignored under certain conditions.

For example, suppose you have three destinations for a domain: A, B, and C. You assign the following preferences:

- A, B, and C at preference 0
- A, B, and C at preference 1
- A, B, and C at preference 2

A, B, and C are all considered reachable.

If a PPP user tries to connect to the domain, suppose the router randomly selects destination A from preference 0. If this connection attempt fails, the router excludes destination A for 5 minutes and goes to the next level (preference 1). From here, it randomly selects destination B, one of the two remaining choices. If the second connection attempt also fails, the router excludes destination B, as well as destination A, and attempts to connect to destination C, the only destination available with preference 2. The router has had an opportunity to connect to every destination available for the domain.

Support for multiple destinations affects the procedure for mapping a user domain name to an L2TP tunnel. To learn how to complete this mapping, see ["Mapping a User Domain Name to an L2TP Tunnel Overview" on page 347](#).

- To enable tunnel selection failover between preference levels:

This tunnel selection method is the default method. If you do not set any tunnel selection parameters, the router uses this method.

### Configuring the Failover Within a Preference Level Method

You use the **l2tp fail-over-within-preference** command to enable tunnel selection failover within a preference level. In this selection method, if the router tries to connect to a destination and is unsuccessful, it selects a new destination at the same preference level. If all destinations at a preference level are marked as unreachable, the router does not attempt to connect to a destination at that level. It drops to the next lower preference level to select a destination.

If all destinations at all preference levels are marked as unreachable, the router chooses the destination that failed first and tries to make a connection. If the connection fails, the router rejects the PPP user session without attempting to contact the remote router.

For example, suppose there are four tunnels for a domain: A, B, C, and D. All tunnels are considered reachable, and the preference levels are assigned as follows:

- A and B at preference 0
- C and D at preference 1

When the router attempts to connect to the domain, suppose it randomly selects tunnel B from preference 0. If it fails to connect to tunnel B, the router excludes tunnel B for five minutes and attempts to connect to tunnel A. If this attempt also fails, the router drops to preference 1. Then suppose the router selects tunnel C. If it also fails to connect to tunnel C, the router excludes tunnel C for five minutes and attempts to connect to tunnel D.

- To enable tunnel selection failover within a preference level:

```
host1(config)#l2tp fail-over-within-preference
```

## Configuring the Maximum Sessions per Tunnel

You can configure the maximum number of sessions per tunnel, either through a RADIUS server or the command-line interface. If you set the maximum sessions per tunnel parameter, the router takes the setting into consideration when it selects a tunnel. If a randomly selected tunnel has a current session count equal to its maximum session count, the router does not attempt to contact that tunnel. Instead, it makes an alternate tunnel selection from the set of reachable tunnels at the same preference level. If no additional reachable tunnels exist at the current preference level, the router drops to the next lower preference level to make the next selection. This process is consistent, regardless of which fail-over scheme is currently running on the router. A tunnel without a configured maximum sessions value has no upper limit on the number of sessions it can support.

The router uses a default value of 0 (zero), which allows unlimited sessions in the tunnel.

- To configure the maximum sessions per tunnel.

```
host1(config)#aaa domain-map lacOne
host1(config-domain-map)#tunnel 1
host1(config-domain-map-tunnel)#max-sessions 1500
```

## Configuring the Weighted Load Balancing Method

With the weighted load-balancing method, the router uses the maximum sessions per tunnel to choose among multiple tunnels that share the same preference level.

The weight of a tunnel is proportional to its maximum session limit and the maximum session limits of the other tunnels at the same preference level. The tunnel with the largest maximum session value has the largest weight; the tunnel with the next largest maximum session value has the next largest weight, down to the tunnel with the smallest maximum session value that has the smallest weight. The router uses a round-robin tunnel selection method by default.

- To configure the router to base tunnel selection within a preference level on the maximum sessions per tunnel.

```
host1(config)#l2tp weighted-load-balancing
```



# Configuring an L2TP LNS

An L2TP network server (LNS) is a node that acts as one side of an L2TP tunnel endpoint and is a peer to the LAC. An LNS is the logical termination point of a PPP connection that is being tunneled from the remote system by the LAC. You can configure your E Series router to function as an LNS.

This chapter includes the following topics that provide information for configuring an L2TP LNS on the E Series router:

- [LNS Configuration Prerequisites on page 366](#)
- [Configuring an LNS on page 366](#)
- [Creating an L2TP Destination Profile on page 369](#)
- [Creating an L2TP Host Profile on page 370](#)
- [Configuring the Maximum Number of LNS Sessions on page 371](#)
- [Configuring Groups for LNS Sessions on page 372](#)
- [Configuring a VR in the L2TP Host Profile on page 373](#)
- [Configuring the RADIUS Connect-Info Attribute on the LNS on page 374](#)
- [Overriding LNS Out-of-Resource Result Codes 4 and 5 on page 375](#)
- [Enabling the L2TP Call Disconnect Notification on Local Pool Exhaustion on page 376](#)
- [Selecting Service Modules for LNS Sessions Using MLPPP on page 377](#)
- [Enabling Tunnel Switching on page 379](#)
- [Creating Persistent Tunnels on page 379](#)
- [Testing Tunnel Configuration on page 380](#)
- [Managing L2TP Destinations, Tunnels, and Sessions on page 380](#)
- [Configuring Disconnect Cause Information on page 381](#)
- [Configuring the Receive Window Size on page 382](#)
- [Configuring Peer Resynchronization on page 385](#)
- [Configuring L2TP Tunnel Switch Profiles on page 389](#)
- [Configuring the Transmit Connect Speed Calculation Method on page 394](#)
- [Controlling the Handling of L2TP Hello Packets in Fast Path on page 402](#)

- [PPP Accounting Statistics on page 403](#)
- [Stateful Line Module Switchover for LNS Sessions on page 405](#)

## LNS Configuration Prerequisites

---

Before you begin configuring the router as an LNS, perform the following steps:

1. Create a virtual router.

```
host1(config)#virtual-router west
```

2. Assign a router ID IP address, such as that for a loopback interface, to the virtual router. This address must be reachable by the L2TP peer.

```
host1:west(config)#ip router-id 10.10.45.3
```



**CAUTION:** You must explicitly assign a router ID to a virtual router rather than using a dynamically assigned router ID. A fixed ID is required because every time the ID changes, L2TP must disconnect all existing tunnels and sessions that use the old ID. If you use a dynamically assigned router ID, the value can change without warning, leading to failure of all L2TP tunnels and sessions. Also, the router could dynamically assign a router ID that is not reachable by the L2TP peer, causing a complete failure of L2TP. You must set the router ID even if you specified a source address in the domain map or a local address in the host profile.

---

### Related Documentation

- [virtual-router](#)
- [ip router-id](#)

## Configuring an LNS

---

When you configure an LNS, you can configure it to accept calls from any LAC.



**NOTE:** If there is no explicit LNS configuration on the router, the UDP port used for L2TP traffic is closed, and no tunnels or sessions can be established.

---

To enable an LAC to connect to the LNS, you must create the following profiles:

- An L2TP destination profile—Defines the location of each LAC
- An L2TP host profile—Defines the attributes used when communicating with an LAC



**NOTE:** If you remove an L2TP destination profile or modify the L2TP host profile configurations (except the following configurations) mandatory for creating the L2TP tunnel, all the existing tunnels and sessions using the profile are dropped.

- Changing the maximum session value
- Changing the session limit group name
- Disabling the single-shot feature
- Enabling or disabling the failover resynchronization
- Assigning or modifying the virtual router

Any changes made to these non-mandatory L2TP destination profile host configurations are applied only to the new tunnels and sessions and not to the existing tunnels and sessions.



**NOTE:** If you are using shared tunnel-server ports, you must configure the shared tunnel-server ports before you configure Layer 2 Tunneling Protocol (L2TP) network server (LNS) support. You use the `tunnel-server` command in Global Configuration mode to specify the physical location of the shared tunnel-server port that you want to configure.

See *virtual-router* for additional information about the `tunnel-server` command and shared tunnel-server ports.

To configure an LNS, perform the following steps:

1. Create a destination profile that defines the location of the LAC, and access L2TP Destination Profile Configuration mode. See [“Creating an L2TP Destination Profile” on page 369](#).

```
host1:boston(config)#l2tp destination profile boston4 ip address 192.168.76.20
host1:boston(config-l2tp-dest-profile)#
```

2. Define the L2TP host profile and enter L2TP Destination Profile Host Configuration mode. See [“Creating an L2TP Host Profile” on page 370](#).

```
host1:boston(config-l2tp-dest-profile)#remote host default
host1:boston(config-l2tp-dest-profile-host)#
```

3. (Optional) Assign a profile name for a remote host.

```
host1:boston(config-l2tp-dest-profile-host)#profile georgeProfile1
```

4. (Optional) Specify a virtual router name, which is used to map PPP session to the TSMs. For more information about configuring a VR in the remote host, see [“Configuring a VR in the L2TP Host Profile” on page 373](#).

```
host1(config-l2tp-dest-profile-host)#session-virtual-router vr1
```

5. (Optional) Disable the use of proxy LCP when connecting to the selected host.

```
host1(config-l2tp-dest-profile-host)#disable proxy lcp
```

6. (Optional) Enable the use of proxy authentication when connecting to the selected host.

```
host1(config-l2tp-dest-profile-host)#enable proxy authenticate
```

7. (Optional) Specify the local hostname to be used in any hostname AVP sends to the LAC. By default, the router name is used as the local hostname.

```
host1(config-l2tp-dest-profile-host)#local host andy
```

8. (Optional) Specify the local IP address to be used in any packets sent to the LAC. By default, the router ID is used.

```
host1(config-l2tp-dest-profile-host)#local ip address 192.168.23.1
```

9. (Optional) Specify the shared secret used to authenticate the tunnel. By default, there is no tunnel authentication.

```
host1:boston(config-l2tp-dest-profile-host)#tunnel password sacco
```

10. (Optional) Specify that the LNS override out-of-resource result codes 4 and 5 with code 2 for interoperation with third-party implementations that do not support codes 4 and 5.

```
host1:boston(config-l2tp-dest-profile-host)#session-out-of-resource-result-code-override
```

11. (Optional) Specify that L2TP create an MLPPP interface when LCP proxy data is not forwarded from the LAC.

For example, the MLPPP interface is created if the LAC does not send the initial received or last received LCP configuration request. If full LCP proxy data is available, this command is ignored.

```
host1:boston(config-l2tp-dest-profile-host)#default-upper-type mlppp
```



**NOTE:** When acting as the LNS, the E Series router supports dialed number identification service (DNIS). With DNIS, if users have a called number associated with them, the router searches the domain map for the called number. If it finds a match, the router uses the matching domain map entry information to authenticate the user. If the router does not find a match, it searches the domain map using normal processing. See the *Using DNIS* section in “[Mapping a User Domain to a Virtual Router Overview](#)” on [page 7](#).

#### Related Documentation

- [Creating an L2TP Destination Profile on page 369](#)
- [Creating an L2TP Host Profile on page 370](#)
- [Configuring the Maximum Number of LNS Sessions on page 371](#)
- [Configuring the RADIUS Connect-Info Attribute on the LNS on page 374](#)
- [Overriding LNS Out-of-Resource Result Codes 4 and 5 on page 375](#)



- [Selecting Service Modules for LNS Sessions Using MLPPP on page 377](#)
- *bundled-group-id*
- *bundled-group-id-overrides-mlppp-ed*
- *default-upper-type mlppp*
- *disable proxy lcp*
- *enable proxy authenticate*
- *l2tp destination profile*
- *local host*
- *local ip address*
- *max-sessions*
- *radius connect-info-format*
- *remote host*
- *session-out-of-resource-result-code-override*
- *session-virtual-router*
- *tunnel password*

---

## Creating an L2TP Destination Profile

You can use the **l2tp destination profile** command to create the destination profile that defines the location of the LAC and to access L2TP Destination Profile Configuration mode.

If no virtual router is specified with the **l2tp destination profile** command, the current virtual router context is used.

If the destination address is 0.0.0.0, then any LAC that can be reached via the specified virtual router is allowed to access the LNS. If the destination address is nonzero, then it must be a host-specific IP address.

To create a destination profile:

- Issue the **l2tp destination profile** command in Global Configuration mode.

```
host1:boston(config)#l2tp destination profile boston ip address 10.10.76.12
host1:boston(config-l2tp-dest-profile)#
```

Use the **no** version to delete the L2TP destination profile.



**NOTE:** When you change an L2TP destination profile, you must wait for the router to delete all L2TP tunnels associated with the deleted profile before you create the new profile.

If you remove a destination profile, all tunnels and sessions using that profile are dropped.

**Related  
Documentation**

- *L2TP/IPsec Tunnels Overview*
- *Configuring an L2TP Destination Profile to Enable IPsec Support for L2TP Tunnels*
- *Configuring Single-Shot L2TP/IPsec Tunnels*
- [Creating an L2TP Host Profile on page 370](#)
- *l2tp destination profile*

---

## Creating an L2TP Host Profile

Use the **remote host** command to define the L2TP host profile and access L2TP Destination Profile Host Configuration mode.

- Each L2TP destination profile can have multiple L2TP host profiles.
- For an LAC to connect to an LNS, the appropriate L2TP destination profile *must* have at least one L2TP host profile.
- If you specify any name other than *default* for the remote host, then the LAC must supply the specified hostname in order for the tunnel to be set up. The remote hostname is matched against the hostname AVP in the received Start-Control-Connection-Request (SCCRQ).
- The remote hostname can be up to 64 characters (no spaces).
- Example

```
host1:boston(config)#l2tp destination profile boston1 ip address 192.168.76.12
host1:boston(config-l2tp-dest-profile)#remote host default
host1(config-l2tp-dest-profile-host)#
```
- Use the **no** version to remove the L2TP host profile.



**NOTE:** If you modify the L2TP host profile configurations (except the following configurations) mandatory for creating the L2TP tunnel, all the existing tunnels and sessions using the profile are dropped.

- Changing the maximum session value
- Changing the session limit group name
- Disabling the single-shot feature
- Enabling or disabling the failover resynchronization
- Assigning or modifying the virtual router

Any changes made to these non-mandatory L2TP destination profile host configurations are applied only to the new tunnels and sessions and not to the existing tunnels and sessions.

- Related Documentation**
- [Creating an L2TP Destination Profile on page 369](#)
  - *l2tp destination profile*

## Configuring the Maximum Number of LNS Sessions

You can use the **max-sessions** command in both L2TP Destination Profile Configuration mode and L2TP Destination Profile Host Configuration mode to configure the number of sessions allowed by the L2TP network server (LNS).

The LNS uses a two-step process to ensure that the maximum number of allowed sessions is not exceeded. When a session is requested, the LNS first checks the maximum sessions set for the L2TP destination profile. If no limit is set, or if the current count is less than the configured limit, the LNS then performs the same check on the L2TP destination host profile limit. If the current count is also less than the L2TP destination host profile limit, then the new session can be established. If a session request exceeds either of the max-sessions settings, the LNS rejects the session.



**NOTE:** New sessions are rejected once the chassis-wide session limit is exceeded, even if the destination profile or host profile maximum session limit is not exceeded. For information about the maximum number of L2TP sessions supported per chassis, see *JunosE Release Notes, Appendix A, System Maximums*.

- To set the maximum sessions allowed for the specified destination, use the **max-sessions** command in L2TP Destination Profile Configuration mode:

```
host1(config)#l2tp destination profile westford ip address 10.10.21.2
host1(config-l2tp-destination-profile)#max-sessions 20000
```

- To set the maximum session allowed for the specified host, use the **max-sessions** command in L2TP Destination Profile Host Configuration mode:

```
host1(config-dest-profile)#remote host default
host1(config-l2tp-destination-profile-host)#max-sessions 20000
```



**CAUTION:** When you modify the maximum number of allowed sessions for the specified destination or host, the LNS checks the number of existing sessions. If the number of existing sessions is greater than the newly configured maximum allowed sessions, the LNS updates the maximum session limit with the newly configured value and then terminates all existing sessions. If the number of existing sessions is less than the newly configured maximum allowed sessions, the LNS updates the maximum session limit with the newly configured value.

Related Documentation

- *max-sessions*

## Configuring Groups for LNS Sessions

You can define and configure session limit groups under the L2TP destination profile. Under each destination profile, you can define a maximum of 4096 session limit groups.

The maximum session limit is applied for each of the session limit groups in L2TP Destination Profile Sessions Limit Group Configuration mode.



**NOTE:** The **max-sessions** command is also supported in L2TP Destination Profile Configuration mode and L2TP Destination Profile Host Configuration mode.

When a session is requested, the LNS first checks the maximum sessions set for the L2TP destination profile. If no limit is set, or if the current session count is less than the configured limit, the LNS then performs the same check on the L2TP destination sessions limit profile. If no limit is set, or if the current session limit is less than the configured limit, the LNS then performs the same check on the L2TP destination host profile limit. If no limit is set, or if the current session count is also less than the L2TP destination host profile limit, then the new session can be established. If a session request exceeds any of the maximum sessions settings, the LNS rejects the session.

To set the maximum sessions allowed for a group for the specified destination, use the **max-sessions** command in L2TP Destination Profile Sessions Limit Group Configuration mode. You can configure this as follows:

1. Define an L2TP destination profile.

```
host1(config)#l2tp destination profile abc virtual-router default ip address 10.10.10.1
```

2. Define a session limit group in L2TP Destination Profile Configuration mode.

```
host1(config-l2tp-dest-profile)#sessions-limit-group g1
```

3. Define the maximum number of sessions allowed in the group.

```
host1(config-l2tp-dest-profile-sessions-limit-group)#max-sessions 8000
```

4. To view the output, use the **show l2tp destination profile** command.

```
host1#show l2tp destination profile abc
```

To set the maximum sessions allowed for a group for the specified host, use the **max-sessions** command in L2TP Destination Profile Sessions Limit Group Configuration mode. You can configure this as follows:

1. Configure a remote hostname.

```
host1(config-l2tp-dest-profile)#remote host xyz
```

2. Assign a sessions limit group name for the remote host.

```
host1(config-l2tp-dest-profile-host)#sessions-limit-group g1
```



**NOTE:** Ensure that the group name is already defined under the destination profile.

3. To view the output, use the **show l2tp destination profile** command.

```
host1#show l2tp destination profile abc
```

#### Related Documentation

- [Configuring the Maximum Number of LNS Sessions on page 371](#)
- *max-sessions*
- *sessions-limit-group*

## Configuring a VR in the L2TP Host Profile

You can configure a VR in an L2TP host profile by using the **session-virtual-router** command to map the corresponding PPP session to the TSMs. By default, no VR is configured in the L2TP host profile and all TSMs are available for all VRs. In this default scenario, the LNS maps a PPP session to a TSM selected in a round-robin manner.

Keep the following points in mind for attaining the desired PPP session to TSM mapping:

- VRs returned from the RADIUS server or configured in Domain Map Configuration mode or Profile Configuration mode must be same as that mentioned in the **session-virtual-router** command. A mismatch between these VR configurations results in mapping PPP session to an unintended ISP TSM.
- You must configure the same VR for the IPv4 and IPv6 subscribers of the same ISP.
- Remote hostnames must be unique for an ISP.

- You can map multiple TSMs to one VR, but you cannot map multiple VRs to one TSM.
- If the VR is mapped to a TSM but not configured in the L2TP host profile, PPP sessions of the VR is handled by the available TSMs that are not mapped to any VR.



---

**NOTE:**

- You must map a VR to TSMs by using the ***l2tp reserve-vr*** command to reserve the TSMs to be used for handling the PPP session managed by the VR. If the VR is not mapped to a TSM, the sessions managed by the VR are handled by the available TSMs that are not mapped to any VR.
  - The PPP session must reside in the configured VR.
- 

To configure a VR in the L2TP host profile:

- Execute the **session-virtual-router** command in L2TP Destination Profile Host Configuration mode.

```
host1(config-l2tp-dest-profile-host)#session-virtual-router vr1
```

You can specify a VR name containing up to 32 alphanumeric characters.

Use the **no** version to remove the configured VR from the L2TP host profile.

If you modify the VR assigned to the L2TP host profile, the TSM manages the already-mapped PPP sessions till the subscriber logs out.

**Related Documentation**

- [Mapping the TSM to a VR](#)
- [Configuring an LNS on page 366](#)
- [Monitoring Configured L2TP Destination Profiles or Host Profiles on page 438](#)
- [L2TP Overview on page 321](#)

---

## Configuring the RADIUS Connect-Info Attribute on the LNS

---

You can configure the LNS to generate the RADIUS Connect-Info attribute [77]. Service providers can then use the information in the RADIUS attribute to identify a customer's service.

On the LNS, the Connect-Info attribute is based on the L2TP connect-speed AVPs received from the LAC. The LNS does not generate the attribute by default. The format of the Connect-Info attribute is as follows, where the TX speed and RX speed are equal to the respective L2TP AVPs:

```
tx-speed [/rx-speed]
```

The TX speed is always included in the attribute when the speed is not zero; however, inclusion of the RX speed depends on the keyword you use with the command.

- Use the **l2tp-connect-speed** keyword to specify that the RX speed is only included when it is not zero and also is different than the TX speed.

```
host1(config)#radius connect-info-format l2tp-connect-speed
```

- Use the **l2tp-connect-speed-rx-when-equal** keyword to specify that the RX speed is always included when it is not zero.

```
host1(config)#radius connect-info-format l2tp-connect-speed-rx-when-equal
```

#### Related Documentation

- [radius connect-info-format](#)

## Overriding LNS Out-of-Resource Result Codes 4 and 5

When the number of L2TP sessions reaches the configured maximum value, the LNS sends an out-of-resource result code (4 or 5) in a CDN (Call-Disconnect-Notify) message to the LAC. This signals the LAC to fail over to another LNS that has the resources for more sessions.

Some third-party LAC implementations fail over only when they receive result code 2 sent in the CDN from the LNS. You can override result codes 4 and 5 with result code 2 on the LNS to enable such routers to fail over to another LNS. These codes have the following meanings:

- 2—Call disconnected for the reason indicated in error code
- 4—Call failed due to lack of appropriate facilities being available (temporary condition)
- 5—Call failed due to lack of appropriate facilities being available (permanent condition)

The following sections describe how to override the result codes and how to display the current code values.

- [Overriding the Result Codes on page 375](#)
- [Displaying the Current Override Setting on page 375](#)

### Overriding the Result Codes

You can override the out-of-resource result codes 4 and 5 by issuing the **session-out-of-resource-result-code-override** command on the LNS.

- To override result codes 4 and 5:

```
host1:boston(config-l2tp-dest-profile-host)#session-out-of-resource-result-code-override
```

### Displaying the Current Override Setting

You can view the current override setting for the LNS result codes in the L2TP destination profile.

- To display the current override setting:

```
ERX(config)#show l2tp destination profile boston
L2TP destination profile boston
Configuration
 Destination address
 Transport ipUdp
```

```
Virtual router default
Peer address 10.10.76.12
Statistics
Destination profile current session count is 0
Host profile attributes
Remote host is LAC
Configuration
Tunnel password is TunnelPass
Local host name is LNS
Local ip address is 46.1.1.2
Session virtual router is vr1
Proxy authenticate is enabled
Disconnect-cause avp is enabled
Tunnels are single-shot
Override out-of-resource-result-code is enabled
Statistics
Current session count is 0
1 L2TP host profile found
```

**Related Documentation**

- *session-out-of-resource-result-code-override*
- *show l2tp destination profile*

---

## Enabling the L2TP Call Disconnect Notification on Local Pool Exhaustion

You can configure the E Series router to enable the LNS to send an L2TP Call-Disconnect-Notify (CDN) message to LAC when an authentication fails because of IPv4 local pool exhaustion in the LNS. Enabling this feature avoids termination of the subscriber only if the LAC is able to tunnel the subscriber to a different LNS on receiving the CDN message. The L2TP CDN message contains the following information:

- Result code set as 2 denoting that the call disconnect reason is indicated in the error code.
- Error code set as 4 denoting that the resources are insufficient to handle the operation now.
- Error message as "session down due to address allocation failure".

If this feature is disabled, then the LNS sends a PPP authentication failure message to the subscriber when the local pool exhausts. The authentication failure message is sent when the address allocation fails due to the following reasons during the address allocation process:

- Address pool is fully utilized.
- Address of the recently logged out subscriber is not yet released to the pool due to the delay timer in LAS.





**NOTE:** The L2TP call disconnect notification support is not applicable for IPv6 pool exhaustion and EAP-based authentication mechanism. In a dual-stack subscriber case, the LNS does not send the CDN message when the IPv6 prefix is assigned to the subscriber but the IPv4 address is not assigned because of IPv4 local pool exhaustion.

To enable the LNS to send an L2TP CDN message to LAC on IPv4 local pool exhaustion:

- Issue the *l2tp disconnect-on-local-pool-exhaustion* command in Global Configuration mode.

```
host1(config)#l2tp disconnect-on-local-pool-exhaustion
```

Use the **no** version to restore the default behavior of sending authentication failure message to the LAC.

#### Related Documentation

- [LNS Configuration Prerequisites on page 366](#)
- [Configuring an LNS on page 366](#)
- [Monitoring Global Configuration Status on E Series Routers on page 433](#)
- [L2TP Overview on page 321](#)
- *l2tp disconnect-on-local-pool-exhaustion*

## Selecting Service Modules for LNS Sessions Using MLPPP

You can install multiple service modules in an E Series router deployed as an LNS where the tunnel sessions carry MLPPP. To use an LNS, at least one Service line module (SM), ES2-S1 Service IOA, or a module that supports the use of shared tunnel-server ports must be installed in the E Series router.

The router selects service modules based on the LNS sessions that underlie the PPP link interfaces of an MLPPP bundle, also known as *bundled sessions*. To determine the appropriate SM where it places the first bundled session for an MLPPP bundle, the router uses a load-balancing mechanism. After the router determines the appropriate SM, it places all sessions for the same bundle on the same SM. By default, the router determines *bundled membership* based on the endpoint discriminator generated by the LNS. The LNS generates its own endpoint discriminator for each MLPPP link, which enables the LNS to handle the MLPPP bundle even if the bundle links are tunneled through multiple LACs.



**NOTE:** The CPEs in a network must validate the endpoint discriminators of bundle links only after the authentication and not during the LCP phase for MLPPP over L2TP to work properly.

For example, an ERX1440 Broadband Services Router has service modules installed in slots 4, 9, and 12. Using the load-balancing mechanism, the router determines that the

SM in slot 4 can accommodate the first bundled session for MLPPP bundle A, and places it there. The first bundled session for bundle A has an endpoint discriminator of 5. The router subsequently places all bundled sessions for bundle A (which have an endpoint discriminator of 5) on the SM in slot 4.

When the SM on which the bundled sessions reside has no more space for additional sessions, the router refuses the L2TP session. This can happen even when other service modules installed in the router have available space.

For more information about endpoint discriminators, see the *Configuring Multilink PPP* chapter in the *JunosE Link Layer Configuration Guide*.

## Assigning Bundled Group Identifiers

In some cases, an endpoint discriminator is not available for the LNS to use to identify the links in a bundled session.

This situation might occur when:

- PPP clients provide endpoint discriminators with null values.
- PPP clients do not provide an endpoint discriminator option when negotiating LCP with the LAC.
- The LAC does not include an endpoint discriminator option in the LCP proxy AVPs.

The router places all bundled sessions without endpoint discriminators on the same SM. However, if there are many such bundled sessions, the load-balanced distribution of LNS sessions across the service modules can deteriorate because the router places all bundled sessions on the same SM without evenly distributing the load.

The **bundled-group-id** command enables you to correct this situation by assigning a numeric bundled group identifier for the router to use when the endpoint discriminator is unavailable to identify the bundled membership. The router places bundled sessions with the same bundled group identifier on the same SM in the same way that it does with endpoint discriminators.

The bundled group identifier applies to the entire router; therefore, if you assign the same bundled group identifier for different L2TP destination host profiles, the router places all of the bundled sessions with the same bundled group identifier on the same SM.



**NOTE:** We recommend that you assign bundled group identifiers only when you are certain that endpoint discriminators are unavailable to identify bundle membership.

---

- To assign a numeric bundled group identifier:  
`host1:boston(config-l2tp-dest-profile-host)#bundled-group-id 4`

## Overriding All Endpoint Discriminators



**NOTE:** We strongly recommend that you use this feature only with the support of JTAC.

You can also configure the router to ignore the value of all endpoint discriminators when it selects a SM and to use only the bundled group identifier that you assigned by issuing the **bundled-group-overrides-mlppp-ed** command.

Issuing the **bundled-group-id** and **bundled-group-id-overrides-mlppp-ed** commands together forces the router to place the bundled sessions on the same SM when a PPP client incorrectly specifies different endpoint discriminators for links in the same bundle.

- To configure the router to ignore the value of all endpoint discriminators:

```
host1:boston(config-l2tp-dest-profile-host)#bundled-group-id-overrides-mlppp-ed
```

### Related Documentation

- *bundled-group-id*
- *bundled-group-id-overrides-mlppp-ed*

## Enabling Tunnel Switching

L2TP tunnel switching allows you to switch packets between one session terminating at an L2TP LNS and another session originating at an L2TP LAC. What distinguishes a tunnel-switched LAC from a conventional one is that there are two interface columns: one for the incoming session (LNS) and one for the outgoing session (LAC). The router forwards traffic from the incoming session to the outgoing session and vice versa.

You can select tunnel switching on a per-chassis basis. By default, tunnel switching is disabled. This preserves current behavior and prevents inadvertent attempts to switch tunnels.



**NOTE:** Each individual L2TP session involved in tunnel switching is counted toward the maximum number of sessions supported on an E Series router.

- To enable tunnel switching:

```
host1(config)#l2tp tunnel-switching
```

### Related Documentation

- *l2tp tunnel-switching*

## Creating Persistent Tunnels

The E Series router supports persistent tunnels. A persistent tunnel is one that is configured to remain available. Persistent tunnels have only local significance; that is, they apply

only to the end of the tunnel where they are set. If the other end of the tunnel chooses to terminate the tunnel, the tunnel is removed.

- To create a persistent tunnel, you configure an idle-timeout value of zero.

```
host1(config)#l2tp tunnel idle-timeout 0
```

**Related Documentation**

- [l2tp tunnel idle-timeout](#)

---

## Testing Tunnel Configuration

You can use the **l2tp tunnel test** command to force the establishment of a tunnel—this enables you to verify both the tunnel configuration and connectivity.

This command supports tunnel initiation: incoming calls on the LAC; outgoing calls on the LNS. The command does not support tunnel respondent: outgoing calls on the LAC; incoming calls on the LNS.

- To test a tunnel configuration:

```
host1#l2tp tunnel test portland.com gold
```

**Related Documentation**

- [l2tp tunnel test](#)

---

## Managing L2TP Destinations, Tunnels, and Sessions

When the router is established as an LNS you can manage the destinations, tunnels and sessions.

- Enable the verification of data integrity via UDP.
- Specify the time period for which the router maintains dynamic destinations, tunnels, or sessions after termination.
- Prevent the creation of new sessions, tunnels, and destinations.
- Close and reopen all or selected destinations, tunnels, and sessions.
- Configure drain timeout operations, which control the amount of time a disconnected LAC tunnel waits before restarting after receiving a restart request.
- Configure how many times the router retries a transmission if the initial attempt is unsuccessful.

**Related Documentation**

- [Generating UDP Checksums in Packets to L2TP Peers on page 333](#)
- [Specifying a Destruct Timeout for L2TP Tunnels and Sessions on page 334](#)
- [Preventing Creation of New Destinations, Tunnels, and Sessions on page 334](#)
- [Shutting Down Destinations, Tunnels, and Sessions on page 336](#)
- [Specifying the Number of Retransmission Attempts on page 337](#)

## Configuring Disconnect Cause Information

You can configure an E Series LNS to convey PPP-related disconnect cause information to its L2TP peer. Enabling an LNS to send disconnect cause information to an LAC is particularly useful in an environment where the LAC initiates tunnels without a client's request, knowledge, or approval. In this type of environment, all PPP signaling for the tunnel session takes place between the LNS and the client, without active participation of the LAC. As a result, the LAC is not aware of the reason that a session has disconnected.



**NOTE:** An E Series LAC does not send PPP Disconnect Case Code AVPs to an LNS. In the event that a third-party LAC does send the AVP to an E Series LNS, the LNS discards the AVP.

1. [Generating the Disconnect Cause AVP Globally on page 381](#)
2. [Generating the Disconnect Cause AVP with a Host Profile on page 381](#)
3. [Enabling RADIUS Accounting for Disconnect Cause on page 382](#)
4. [Displaying Disconnect Cause Statistics on page 382](#)

### Generating the Disconnect Cause AVP Globally

You use the **l2tp disconnect-cause** command to specify that the LNS include the PPP Disconnect Cause Code AVP in all L2TP Call-Disconnect-Notify (CDN) messages that it sends to the LAC. For example, this feature enables the LAC to obtain information about the cause of a session disconnection,

- To enable disconnect cause generation chassis-wide on the LNS:

```
host1(config)#l2tp disconnect-cause
```



**NOTE:** Sessions for which the AVP generation is enabled by the host-profile-specific **disconnect-cause** command continue to generate the AVP.

### Generating the Disconnect Cause AVP with a Host Profile

You use the **disconnect-cause** command in L2TP Destination Profile Host Configuration mode to specify that the E Series LNS generate PPP Disconnect Cause Code AVPs. This command pertains only to L2TP sessions to which the L2TP destination host profile applies. The AVP is included in all L2TP CDN messages that the LNS sends to an LAC for covered sessions.



**NOTE:** This command is used only for dial-in sessions; use the **l2tp disconnect-cause** command in Global Configuration mode to generate PPP Disconnect Cause Code AVPs for dial-out sessions.

- To enable disconnect cause generation for all tunnels that use a particular host profile on the LNS:

```
host1(config-l2tp-dest-profile-host)#disconnect-cause
```

## Enabling RADIUS Accounting for Disconnect Cause

You use the **radius include l2tp-ppp-disconnect-cause acct-stop enable** command to specify that the Disconnect-Cause RADIUS attribute (VSA 26-51) is generated and included in RADIUS acct-stop and acct-tunnel-link-stop records. RADIUS VSA 26-51 is not included in the accounting records by default.

At the LAC, this accounting reports remotely generated disconnect cause information received from the LNS. At the LNS, the accounting reports locally generated disconnect cause information.

- To enable disconnect cause accounting:

```
host1(config)#radius include l2tp-ppp-disconnect-cause acct-stop enable
```

## Displaying Disconnect Cause Statistics

You can display chassis-wide summary statistics for all disconnect cause information received by the LAC, sorted by code number.

- To display summary statistics for disconnect cause information:

```
host1(config)#show l2tp received-disconnect-cause-summary
```

## Configuring the Receive Window Size

---

You can configure the L2TP receive window size (RWS) for an L2TP tunnel. L2TP uses the RWS to implement a sliding window mechanism for the transmission of control messages.

When you configure the RWS, you specify the number of packets that the L2TP peer can transmit without receiving an acknowledgment from the router. If the RWS is not configured, the router determines the RWS value (always set as 64) and uses this value for all new tunnels on both the LAC and the LNS.

You can configure the L2TP RWS in the following ways:

- Configure the systemwide default RWS setting for a tunnel on both the LAC and the LNS by using the **l2tp tunnel default-receive-window** command (in global Configuration mode).
- Configure the RWS for a tunnel on the LAC by using either the **receive-window** command (in Domain Map Tunnel Configuration mode) or by including the L2tp-Recv-Window-Size RADIUS attribute (VSA 26-54) in RADIUS Access-Accept messages.

- Configure the RWS for all tunnels that use a particular host profile on the LNS by using the **receive-window** command (in L2TP Destination Profile Host Configuration mode).

1. [Configuring the Default Receive Window Size on page 383](#)
2. [Configuring the Receive Window Size on the LAC on page 383](#)
3. [Configuring the Receive Window Size on the LNS on page 384](#)

## Configuring the Default Receive Window Size

Use the **l2tp tunnel default-receive-window** command to configure the default L2TP RWS for a tunnel on both the LAC and the LNS. The default L2TP RWS is the number of packets that the L2TP peer can transmit without receiving an acknowledgment from the router. The only supported value is 4.

To configure the default RWS setting:

1. From Global Configuration mode, set the L2TP default RWS. The only value supported for the default RWS is 4.

```
host1(config)#l2tp tunnel default-receive-window 4
```

The router uses this RWS value for all new tunnels on both the LAC and the LNS. The new command has no effect on previously configured tunnels.

2. (Optional) Use the **show l2tp** command to verify the default RWS configuration.

```
host1#show l2tp
Configuration
 L2TP administrative state is enabled
 Dynamic interface destruct timeout is 600 seconds
 Data packet checksums are disabled
 Receive data sequencing is not ignored
 Tunnel switching is disabled
 Retransmission retries for established tunnels is 5
 Retransmission retries for not-established tunnels is 5
 Tunnel idle timeout is 60 seconds
 Failover within a preference level is disabled
 Weighted load balancing is disabled
 Tunnel authentication challenge is enabled
 Calling number avp is enabled
 Ignore remote transmit address change is disabled
 Disconnect cause avp is disabled
 Default receive window size is 4
Sub-interfaces total active failed auth-errors
Destinations 0 0 0 n/a
Tunnels 0 0 0 0
Sessions 0 0 0 n/a
Switched-sessions 0 0 0 n/a
```

## Configuring the Receive Window Size on the LAC

Use the **receive-window** command to configure the L2TP RWS for a tunnel on the LAC. Use the **no** version of the command to revert to the systemwide RWS setting configured with the **l2tp tunnel default-receive-window** command.



**TIP:** The RWS setting must be the same for all users of the same tunnel.

If you modify the RWS setting for an existing tunnel, subsequent tunnel users might not be able to log in if their RWS setting conflicts with the new RWS setting for the tunnel.

To configure the RWS for a tunnel on the LAC:

1. Access Domain Map Tunnel Configuration mode as described in [“Mapping a User Domain Name to an L2TP Tunnel Overview” on page 347](#). For example:

```
host1(config)#aaa domain-map fms.com
host1(config-domain-map)#router-name westford
host1(config-domain-map)#tunnel 3
host1(config-domain-map-tunnel)#
```

2. From Domain Map Tunnel Configuration mode, set the tunnel RWS. The only value supported for the tunnel RWS is 4, and it must be the same for all users of the same tunnel.

```
host1(config-domain-map-tunnel)#receive-window 4
```

3. (Optional) Use the **show aaa domain-map** command to verify the RWS configuration.

```
host1#show aaa domain-map
```

```
Domain: fms.com; router-name: westford; ipv6-router-name: default
```

| Tunnel<br>Tag | Tunnel<br>Peer           | Tunnel<br>Source     | Tunnel<br>Type            | Tunnel<br>Medium | Tunnel<br>Password | Tunnel<br>Id | Tunnel<br>Client<br>Name |
|---------------|--------------------------|----------------------|---------------------------|------------------|--------------------|--------------|--------------------------|
| 3             | <null>                   | <null>               | l2tp                      | ipv4             | <null>             | <null>       | <null>                   |
| Tunnel<br>Tag | Tunnel<br>Server<br>Name | Tunnel<br>Preference | Tunnel<br>Max<br>Sessions | Tunnel<br>RWS    |                    |              |                          |
| 3             | <null>                   | 2000                 | 0                         | 4                |                    |              |                          |

You can also configure the RWS for a tunnel on the LAC by including the L2tp-Recv-Window-Size RADIUS attribute (VSA 26-54) in RADIUS Access-Accept messages. For more information about RADIUS Access-Accept messages, see [“Subscriber AAA Access Messages Overview” on page 182](#). For more information about the L2tp-Recv-Window-Size attribute, see [“RADIUS IETF Attributes” on page 241](#).

## Configuring the Receive Window Size on the LNS

Use the **receive-window** command to configure the L2TP RWS for a tunnel on the LNS. Use the **no** version of the command to revert to the systemwide RWS setting configured with the **l2tp tunnel default-receive-window** command.

To configure the RWS for a tunnel on the LNS:

1. Access L2TP Destination Profile Host Configuration mode. For example:

```
host1(config)#virtual-router fms02
host1:fms02(config)#l2tp destination profile fms02 ip address 192.168.5.61
```



```
host1:fms02(config-l2tp-dest-profile)#remote host fms03
host1:fms02(config-l2tp-dest-profile-host)#
```

2. From Destination Profile Host Configuration mode, set the tunnel RWS. The only value supported for the tunnel RWS is 4.

```
host1:fms02(config-l2tp-dest-profile-host)#receive-window 4
```



**TIP:** If you modify the RWS setting of a host profile for an existing tunnel, the router drops the tunnel. This action is consistent with router behavior when you modify an L2TP host profile.

3. (Optional) Use the ***show l2tp destination profile*** command to verify the RWS configuration.

```
host1:fms02#show l2tp destination profile fms02
L2TP destination profile fms02
Destination address
 Transport ipUdp
 Virtual router fms02
 Peer address 192.168.5.61
Host profile attributes
 Remote host is fms03
 Receive window size is 4
.
.
.
1 L2TP host profile found
```

#### Related Documentation

- [Monitoring Mapping Between User Domains and Virtual Routers on page 120](#)
- [Monitoring the Mapping for User Domains and Virtual Routers with AAA on page 427](#)
- [Monitoring Configured Tunnel Groups with AAA on page 430](#)
- [Monitoring Global Configuration Status on E Series Routers on page 433](#)
- [Monitoring Configured L2TP Destination Profiles or Host Profiles on page 438](#)

## Configuring Peer Resynchronization

The JunosE Software enables you to configure the peer resynchronization method you want the router to use. Peer resynchronization enables L2TP to recover from a router warm start and to allow an L2TP failed endpoint to resynchronize with its peer non-failed endpoint.

L2TP peer resynchronization:

- Prevents the non-failed endpoint from prematurely terminating a tunnel while the failed endpoint is recovering
- Reestablishes the sequence numbers required for the operation of the L2TP control protocol

- Resolves inconsistencies in the tunnel and session databases of the failed endpoint and the non-failed endpoint

To ensure successful peer resynchronization between endpoints, the non-failed endpoint must support a complete RFC-compliant L2TP implementation.

JunosE Software supports both the L2TP silent failover method and the L2TP failover protocol method, which is described in Fail Over extensions for L2TP “failover” draft-ietf-l2tpext-failover-06.txt. You can configure L2TP to use the failover protocol method as the primary peer resynchronization method, but then fall back to the silent failover method if the peer does not support the failover protocol method.

The following list highlights differences between the failover protocol and silent failover peer resynchronization methods:

- With the L2TP failover protocol method, both endpoints must support the method or recovery always fails. The L2TP failover protocol method also requires a non-failed endpoint to wait an additional recovery time period while the failed endpoint is recovering to prevent the non-failed endpoint from prematurely disconnecting the tunnel. The additional recovery period makes L2TP less responsive to the loss of tunnel connectivity.
- Silent failover operates entirely within the failed endpoint and does not require non-failed endpoint support—this improves interoperability between peers. Silent failover does not require additional recovery time by the non-failed endpoint, which also eliminates the potential for degraded responsiveness to the loss of tunnel connectivity.



**NOTE:** L2TP silent failover is not supported on E3 ATM and CT1 line modules in peer-facing configurations.

---



**NOTE:** If an LNS device at one end of an L2TP tunnel encounters a failure and is not configured with the L2TP peer resynchronization method to enable the LNS device to resynchronize with the non-failed endpoint peer (the LAC device at the other end of the tunnel), the tunnel is brought down immediately after the configured value for the number of retransmission attempts is exceeded. The tunnel between the LAC device and the failed LNS device that is recovering is not preserved for the default recovery time period, which is 15 minutes. Instead, the tunnel is terminated immediately and the LAC device sends the Failover Capability attribute-value pair (AVP) in the Stop-Control-Connection-Notification (StopCCN) packet to the original address with a failover recovery time field set to zero.

---

You can use the CLI or RADIUS to configure the resynchronization method for your router.

1. [Configuring Peer Resynchronization for L2TP Host Profiles and AAA Domain Map Tunnels on page 387](#)
2. [Configuring the Global L2TP Peer Resynchronization Method on page 388](#)
3. [Using RADIUS to Configure Peer Resynchronization on page 389](#)

## Configuring Peer Resynchronization for L2TP Host Profiles and AAA Domain Map Tunnels

The JunosE CLI enables you to configure the peer resynchronization method globally, for a host profile, or for a domain map tunnel. A host profile or domain map tunnel configuration takes precedence over the global peer resynchronization configuration.

When you change the peer resynchronization method, the change is not immediately applied to existing tunnels. Tunnels continue using their current resynchronization method until the next time the tunnel is reestablished.

Use the **failover-resync** command to configure the L2TP peer resynchronization method for L2TP host profiles and AAA domain map tunnels. This command takes precedence over the global peer resynchronization configuration.

Choose one of the following keywords to specify the peer resynchronization method:

- **failover-protocol**—The tunnel uses the L2TP failover protocol method. If the peer non-failed endpoint does not support the L2TP failover protocol, a failover forces disconnection of the tunnel and all of its sessions.
- **failover-protocol-fallback-to-silent-failover**—The tunnel uses the L2TP failover protocol method; however, if the peer non-failed endpoint does not support the L2TP failover protocol method, the tunnel falls back to using the silent failover method.
- **silent-failover**—The tunnel uses the silent failover method. The tunnel also informs its peer that it supports the failover protocol method for the peer's failovers.
- **disable**—The tunnel does not use any peer resynchronization method for its own failovers. The tunnel informs its peer that it supports the failover protocol method for the peer's failovers. A failover forces the disconnection of the tunnel and all of its sessions.
- **not-configured**—Peer resynchronization is not configured for L2TP host profiles and AAA domain map tunnels. L2TP uses the global failover method.

By default, peer resynchronization is not configured at the L2TP profile-level or the domain map-level—therefore, the global configuration is used. This is different than using the **disable** keyword, which specifies that no peer synchronization method is used.

Use the **show l2tp destination profile** command to display a host profile's peer resynchronization configuration and the **show aaa domain-map** command to display a domain map's configuration.

- To configure peer resynchronization for an L2TP host profile:
 

```
host1(config)#l2tp destination profile lac-dest ip address 192.168.20.2
host1(config-l2tp-dest-profile)#remote host lac-host
```

```
host1(config-l2tp-dest-host-profile-host)#failover-resync silent-failover
```

- To configure peer resynchronization for an AAA domain map tunnel:

```
host1(config)#aaa domain-map lac-tunnel
host1(config-domain-map)#tunnel 10
host1(config-domain-map-tunnel)#failover-resync silent-failover
```

## Configuring the Global L2TP Peer Resynchronization Method

You can configure the peer resynchronization method globally, or for L2TP host profiles or domain map tunnels—a host profile or domain map tunnel configuration takes precedence over the global peer resynchronization configuration.

When you change the peer resynchronization method, the change is not immediately applied to existing tunnels. Tunnels continue using their current resynchronization method until the next time the tunnel is reestablished.

Use the **l2tp failover-resync** command to configure the global L2TP peer resynchronization method that L2TP failed endpoints use to resynchronize with a peer non-failed endpoint.

Choose one of the following keywords to specify the peer resynchronization method. All tunnels in the chassis use the specified method unless it is overridden by an L2TP host profile configuration or an AAA domain map configuration.

- **failover-protocol**—Tunnels use the L2TP failover protocol method. If the peer non-failed endpoint does not support the L2TP failover protocol, a failover forces disconnection of all tunnels and their sessions.
- **failover-protocol-fallback-to-silent-failover**—Tunnels use the L2TP failover protocol method; however, if the peer non-failed endpoint does not support the L2TP failover protocol method, the tunnel falls back to using the silent failover method.
- **silent-failover**—Tunnels use the silent failover method. The tunnels also inform their peers that they support the failover protocol method for peer failovers.
- **disable**—Tunnels do not use any peer resynchronization method for their own failovers. Tunnels inform their peers that they support the failover protocol method for peer failovers. A failover forces the disconnection of all tunnels and sessions.

Use the **show l2tp** command to display the global peer resynchronization configuration.

- To configure peer resynchronization for an L2TP host profile or AAA domain map tunnel:

```
host1(config)#l2tp failover-resync silent-failover
```

- To restore the global default setting, which uses the **failover-protocol-fallback-to-silent-failover** method:

```
host1(config)#default l2tp failover-resync
```

- To disable peer resynchronization, use the **no** version of the command—this is the same as using the **disable** keyword:

```
host1(config)#no l2tp failover-resync
```

## Using RADIUS to Configure Peer Resynchronization

The JunosE Software supports the use of RADIUS to configure the L2TP peer resynchronization method used by your L2TP tunnels. You use the L2TP-Resynch-Method RADIUS attribute (VSA 26-90) in RADIUS Access-Accept messages to specify the L2TP peer resynchronization method.

[Table 78 on page 389](#) describes the L2TP-Resynch-Method RADIUS attribute. For more information about RADIUS Access-Accept messages, see [“Subscriber AAA Access Messages Overview” on page 182](#). For more information about the L2TP-Resynch-Method attribute, see [“RADIUS IETF Attributes” on page 241](#).

**Table 78: L2TP-Resynch-Method RADIUS Attribute**

| Standard Number | Attribute Name      | Description                        | Length | Subtype Length | Value                                                                                                                                                                                                    |
|-----------------|---------------------|------------------------------------|--------|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| [26-90]         | L2TP-Resynch-Method | L2TP peer resynchronization method | 12     | 6              | integer: <ul style="list-style-type: none"> <li>• 0 = disabled</li> <li>• 1= failover protocol</li> <li>• 2 = silent failover</li> <li>• 3 = failover protocol with silent failover as backup</li> </ul> |

## Configuring L2TP Tunnel Switch Profiles

You can use the **l2tp switch-profile** command to create an L2TP tunnel switch profile. An *L2TP tunnel switch profile* is a set of characteristics that defines the behavior of L2TP tunnel switching for the interfaces to which the profile is assigned.

Within the L2TP tunnel switch profile, you configure a particular tunnel switching behavior for a specified L2TP AVP. For example, you can configure the router to preserve the value of (relay) a specified AVP type across the LNS/LAC boundary in an L2TP tunnel-switched network.

## Applying the L2TP Tunnel Switch Profile

Configuring an L2TP tunnel switch profile has no effect by itself. To use the tunnel switch profile in an L2TP tunnel-switched network, you must apply it to an L2TP outbound LAC session by using one of the following methods:

- Authentication, authorization, and accounting (AAA) domain maps
- AAA tunnel groups
- RADIUS Access-Accept messages

If none of these methods are used, you can apply the L2TP tunnel switch profile as an AAA default tunnel parameter. The default tunnel switch profile has lower precedence than the other methods for applying the tunnel switch profile.

For more information about the methods for applying L2TP tunnel switch profiles, see [“Configuration Tasks” on page 390](#).

## Configuration Guidelines

The following rules apply when you configure L2TP tunnel switch profiles:

- L2TP tunnel switching must be enabled for tunnel switch profiles to take effect. For information, see [“Enabling Tunnel Switching” on page 379](#).
- L2TP tunnel switch profiles have no effect when they are assigned to a LAC session that is not tunnel switched.
- The router can relay only those AVPs that are accepted at the LNS. Malformed AVPs are never relayed.
- If a tunnel grant response specifies a named tunnel switch profile that has not been configured on the router, the router prohibits connection of the L2TP tunnel-switched session.
- If you remove a tunnel switch profile, the router also disconnects all associated L2TP switched sessions using that profile.
- In some cases, attributes configured in a tunnel switch profile take precedence over similar attributes configured globally on the router.

For example, configuring L2TP Calling Number AVP 22 for relay overrides the **l2tp disable calling-number-avp** command issued from Global Configuration mode to prevent the router from sending AVP 22 in incoming-call-request (ICRQ) packets. In this scenario, the router relays the Calling Number AVP.

## Configuring L2TP AVPs for Relay

Previously, the router did not preserve the values of incoming L2TP AVPs across the LNS/LAC boundary in an L2TP tunnel-switched network. The router regenerated most incoming AVPs, such as L2TP Calling Number AVP 22, based on the local policy in effect. However, some AVPs, such as Cisco NAS Port Info AVP 100, were dropped.

In an L2TP tunnel switch profile, you can define the types of AVPs that the router can relay unchanged across the LNS/LAC boundary. You can specify that the router relay one or more of the following AVP types:

- L2TP Bearer Type AVP 18
- L2TP Calling Number AVP 22
- Cisco NAS Port Info AVP 100

When you configure any of these AVP types for relay in an L2TP tunnel-switched network, the router preserves the value of an incoming AVP of this type when packets are switched between the inbound LNS session and the outbound LAC session.

## Configuration Tasks

To configure and use an L2TP tunnel switch profile in an L2TP tunnel-switched network:

1. Ensure that L2TP tunnel switching is enabled on the router.
2. Configure the L2TP tunnel switch profile.
3. Apply the L2TP tunnel switch profile to the tunnel in one of the following ways:
  - To apply a named tunnel switch profile through an AAA domain map, use the **switch-profile** command from Domain Map Tunnel Configuration mode. For details, see [“Applying L2TP Tunnel Switch Profiles by Using AAA Domain Maps” on page 392](#).
  - To apply a named tunnel switch profile through an AAA tunnel group, use the **switch-profile** command from Tunnel Group Tunnel Configuration mode. For details, see [“Applying L2TP Tunnel Switch Profiles by Using AAA Tunnel Groups” on page 393](#).
  - To apply a named tunnel switch profile through RADIUS, include the Tunnel-Switch-Profile RADIUS attribute (VSA 26-91) in RADIUS Access-Accept messages. For details, see [“Applying L2TP Tunnel Switch Profiles by Using RADIUS” on page 394](#).
  - To apply a default tunnel switch profile to a virtual router, use the **aaa tunnel switch-profile** command from Global Configuration mode. For details, see [“Applying Default L2TP Tunnel Switch Profiles” on page 393](#).

The following sections describe how to perform each of these tasks.

### Enabling Tunnel Switching on the Router

To enable L2TP tunnel switching on the router, use the **l2tp tunnel-switching** command. By default, tunnel switching is disabled.

- To enable L2TP tunnel switching:
 

```
host1(config)#l2tp tunnel-switching
```

For more information, see [“Enabling Tunnel Switching” on page 379](#).

### Configuring L2TP Tunnel Switch Profiles

To configure an L2TP tunnel switch profile:

1. Create the L2TP tunnel switch profile and assign it a name. The **l2tp switch-profile** command accesses L2TP Tunnel Switch Profile Configuration mode.
 

```
host1(config)#l2tp switch-profile concord
host1(config-l2tp-tunnel-switch-profile)#
```
2. Configure the L2TP tunnel switching behavior for the interfaces to which this profile is assigned. Use the **avp** command with the **relay** keyword to cause the router to preserve the value of an incoming AVP of this type when packets are switched between an inbound LNS session and an outbound LAC session.

You can use any of the following keywords to specify the AVPs for the router to relay:

- **bearer-type**—L2TP Bearer Type AVP 18; by default, the router regenerates this AVP at the outbound LAC session, based on the local policy in effect
- **calling-number**—L2TP Calling Number AVP 22; by default, the router regenerates this AVP at the outbound LAC session, based on the local policy in effect

- **cisco-nas-port**—Cisco NAS Port Info AVP 100; by default, the router drops this AVP

Use the **no** version to restore the default L2TP tunnel switching behavior (regenerate or drop) for incoming AVPs of the specified type.

The following commands configure the router to relay the Bearer Type, Calling Number, and Cisco NAS Port Info AVP types across the LNS/LAC boundary.

```
host1(config-l2tp-tunnel-switch-profile)#avp bearer-type relay
host1(config-l2tp-tunnel-switch-profile)#avp calling-number relay
host1(config-l2tp-tunnel-switch-profile)#avp cisco-nas-port relay
```

3. (Optional) Use the **show l2tp switch-profile** command to verify configuration of the tunnel switch profile.

```
host1(config-l2tp-tunnel-switch-profile)# run show l2tp switch-profile
L2TP tunnel switch profile concord
L2TP tunnel switch profile myProfile
2 L2TP tunnel switch profiles found
host1(config-l2tp-tunnel-switch-profile)# run show l2tp switch-profile concord
L2TP tunnel switch profile concord
 AVP bearer type action is relay
 AVP calling number action is relay
 AVP Cisco nas port info action is relay
```

### Applying L2TP Tunnel Switch Profiles by Using AAA Domain Maps

To apply an L2TP tunnel switch profile to sessions associated with an AAA domain map:

1. Access Domain Map Tunnel Configuration mode.

```
host1(config)#aaa domain-map westford.com
host1(config-domain-map)#router-name default
host1(config-domain-map)#tunnel 3
host1(config-domain-map-tunnel)#
```

For more information about how to map a domain to an L2TP tunnel from Domain Map Tunnel Configuration mode, see [“Mapping a User Domain Name to an L2TP Tunnel Overview” on page 347](#).

2. From Domain Map Tunnel Configuration mode, issue the **switch-profile** command to apply the specified L2TP switch profile to the sessions associated with this domain map.

```
host1(config-domain-map-tunnel)#switch-profile concord
```

3. (Optional) Use the **show aaa domain-map** command to verify application of the tunnel switch profile.

```
host1(config-domain-map-tunnel)#run show aaa domain-map
Domain: westford.com; router-name: default; ipv6-router-name: default
```

| Tunnel<br>Tag | Tunnel<br>Peer | Tunnel<br>Source | Tunnel<br>Type | Tunnel<br>Medium | Tunnel<br>Password | Tunnel<br>Id | Tunnel<br>Client<br>Name |
|---------------|----------------|------------------|----------------|------------------|--------------------|--------------|--------------------------|
| 3             | <null>         | <null>           | l2tp           | ipv4             | <null>             | <null>       | <null>                   |

| Tunnel<br>Tag | Tunnel<br>Server<br>Name | Tunnel<br>Preference | Tunnel<br>Max<br>Sessions | Tunnel<br>RWS | Tunnel<br>Virtual<br>Router | Tunnel<br>Switch<br>Profile |
|---------------|--------------------------|----------------------|---------------------------|---------------|-----------------------------|-----------------------------|
|               |                          |                      |                           |               |                             |                             |



```

3 <null> 2000 0 system chooses <null> concord

```

### Applying L2TP Tunnel Switch Profiles by Using AAA Tunnel Groups

To apply an L2TP tunnel switch profile to sessions associated with an AAA tunnel group:

1. Access Tunnel Group Tunnel Configuration mode.

```

host1(config)#aaa tunnel-group sunnyvale
host1(config-tunnel-group)#tunnel 3
host1(config-tunnel-group-tunnel)#

```

For more information about how to map a domain to an L2TP tunnel from Tunnel Group Tunnel Configuration mode, see [“Mapping a User Domain Name to an L2TP Tunnel Overview” on page 347](#).

2. From Tunnel Group Tunnel Configuration mode, issue the **switch-profile** command to apply the specified L2TP switch profile to the sessions associated with this tunnel group.

```

host1(config-tunnel-group-tunnel)#switch-profile sanjose

```

3. (Optional) Use the **show aaa tunnel-group** command to verify application of the tunnel switch profile.

```

host1(config-tunnel-group-tunnel)#run show aaa tunnel-group

```

Tunnel Group: sunnyvale

| Tunnel Tag | Tunnel Peer | Tunnel Source     | Tunnel Type         | Tunnel Medium  | Tunnel Password       | Tunnel Id             | Tunnel Client Name |
|------------|-------------|-------------------|---------------------|----------------|-----------------------|-----------------------|--------------------|
| 3          | <null>      | <null>            | l2tp                | ipv4           | <null>                | <null>                | <null>             |
| Tunnel Tag | Server Name | Tunnel Preference | Tunnel Max Sessions | Tunnel RWS     | Tunnel Virtual Router | Tunnel Switch Profile |                    |
| 3          | <null>      | 2000              | 0                   | system chooses | <null>                | sanjose               |                    |

### Applying Default L2TP Tunnel Switch Profiles

You can apply a default L2TP tunnel switch profile to a virtual router by issuing the **aaa tunnel switch-profile** command from Global Configuration mode. The router uses the default tunnel switch profile if the tunnel attributes returned from an AAA domain map or tunnel group or from a RADIUS authentication server *do not include* a named tunnel switch profile. The router ignores the default tunnel switch profile if the tunnel attributes returned from an AAA domain map or tunnel group or from a RADIUS authentication server *do include* a named tunnel switch profile.

The default L2TP tunnel switch profile applies to a specific virtual router. You can apply a different default tunnel switch profile to each virtual router configured.

To apply a default L2TP tunnel switch profile to a virtual router:

1. Create the virtual router to which you want to apply the default tunnel switch profile.

```

host1(config)#virtual-router east

```

```
host1:east(config)#
```

2. Issue the **aaa tunnel switch-profile** command to apply the default L2TP tunnel switch profile in the context of this virtual router.

```
host1:east(config)#aaa tunnel switch-profile boston
```

3. (Optional) Use the **show aaa tunnel-parameters** command to verify application of the default tunnel switch profile.

```
host1:east(config)#run show aaa tunnel-parameters
Tunnel password is <NULL>
Tunnel client-name is <NULL>
Tunnel nas-port-method is none
Tunnel switch-profile is boston
Tunnel nas-port ignore disabled
Tunnel nas-port-type ignore disabled
Tunnel assignmentId format is assignmentId
Tunnel calling number format is descriptive
```

---

### Applying L2TP Tunnel Switch Profiles by Using RADIUS

On the LAC, the router can receive tunnel configuration attributes through a RADIUS authentication server. To use RADIUS to apply an L2TP tunnel switch profile to a session, you can configure RADIUS to include the Tunnel-Switch-Profile RADIUS attribute (VSA 26-91) in RADIUS Access-Accept messages.

For more information about RADIUS Access-Accept messages, see [“Subscriber AAA Access Messages Overview” on page 182](#). For more information about the Tunnel-Switch-Profile attribute, see [“RADIUS IETF Attributes” on page 241](#).

#### Related Documentation

- [Enabling Tunnel Switching on the Router on page 391](#)
- [Configuring L2TP Tunnel Switch Profiles on page 391](#)
- [Applying L2TP Tunnel Switch Profiles by Using AAA Domain Maps on page 392](#)
- [Applying L2TP Tunnel Switch Profiles by Using AAA Tunnel Groups on page 393](#)
- [Applying Default L2TP Tunnel Switch Profiles on page 393](#)
- [Applying L2TP Tunnel Switch Profiles by Using RADIUS on page 394](#)
- *aaa tunnel switch-profile*
- *avp*
- *l2tp switch-profile*
- *l2tp tunnel-switching*

---

### Configuring the Transmit Connect Speed Calculation Method

You can configure the method that the router uses to calculate the transmit connect speed of the subscriber's access interface for a tunneled L2TP session. L2TP reports the transmit connect speed in L2TP Transmit (TX) Speed AVP 24. During the establishment of an L2TP tunnel session, the LAC sends AVP 24 to the LNS to convey the transmit speed of the subscriber's access interface.

You can configure the calculation method for the transmit connect speed reported in L2TP Transmit (TX) Speed AVP 24 in any of the following ways. The first three methods—AAA domain maps, AAA tunnel groups, and RADIUS—are mutually exclusive.

- AAA domain maps—Use the **tx-connect-speed-method** command from Domain Map Tunnel Configuration mode. For instructions, see [“Using AAA Domain Maps to Configure the Transmit Connect Speed Calculation Method” on page 399](#).
- AAA tunnel groups—Use the **tx-connect-speed-method** command from Tunnel Group Tunnel Configuration mode. For instructions, see [“Using AAA Tunnel Groups to Configure the Transmit Connect Speed Calculation Method” on page 399](#).
- AAA default tunnel parameters—Use the **aaa tunnel tx-connect-speed-method** command from Global Configuration mode. The router uses the calculation method specified with this command if the tunnel attributes returned from an AAA domain map, an AAA tunnel group, or a RADIUS authentication server do not include the transmit connect speed calculation method. For instructions, see [“Using AAA Default Tunnel Parameters to Configure the Transmit Connect Speed Calculation Method” on page 400](#).
- RADIUS Include the Tunnel-Tx-Speed-Method RADIUS attribute (Juniper Networks VSA 26-94) in RADIUS Access-Accept messages. For instructions, see [“Using AAA Default Tunnel Parameters to Configure the Transmit Connect Speed Calculation Method” on page 400](#).

## Transmit Connect Speed Calculation Methods

In previous releases, the router calculated the transmit speed of the subscriber's access interface based only on statically configured settings for the underlying layer 2 access interface. With this feature, you can obtain a more accurate representation of the transmit connect speed by choosing a calculation method that reflects changes to the layer 2 interface due to statically configured settings, dynamically configured settings, or QoS settings.

You can choose one of the following methods for calculating the transmit connect speed that is reported in L2TP Transmit (TX) Speed AVP 24:

- Static layer 2
- Dynamic layer 2
- QoS
- Actual (lesser of dynamic layer 2 or QoS)

The following sections describe each of these calculation methods.



**NOTE:** Configuring the transmit connect speed calculation method has no effect on the operation of the L2TP Receive (RX) Speed AVP 38 or the Connect-Info RADIUS attribute [77] at the LAC.

## Static Layer 2

---

The static layer 2 method calculates the transmit connect speed of the subscriber's access interface based on the statically configured settings for the underlying layer 2 ATM 1483 or Ethernet interface. The static layer 2 method does not reflect changes to the transmit speed of the layer 2 interface due to dynamically configured settings or to QoS.

For ATM 1483 circuits, the static layer 2 value is based on the bandwidth that the connection requires. The router uses certain traffic parameters for each service category to determine the required bandwidth for the connection. For more information about how the router computes bandwidth for ATM 1483 circuits, see the *Connection Admission Control* section in the *JunosE Link Layer Configuration Guide*.

For Ethernet VLANs, the static layer 2 value is the advisory transmit speed of the VLAN subinterface, if configured with the **vlan advisory-tx-speed** command, or the speed of the underlying physical port if the advisory transmit speed is not configured.

If there is no explicit static configuration for the layer 2 interface, L2TP reports the speed of the underlying physical port as the transmit connect speed.

## Dynamic Layer 2

---

The dynamic layer 2 method calculates the transmit connect speed of the subscriber's access interface based on the dynamically configured settings for the underlying layer 2 interface.

If there is no dynamic configuration for the layer 2 interface, L2TP reports the transmit connect speed based on statically configured settings. If there is no static speed configuration for the layer 2 interface, L2TP reports the speed of the underlying physical port as the transmit connect speed.

## QoS

---

The QoS method calculates the transmit connect speed of the subscriber's access interface based on settings determined by static or dynamic QoS configurations. This calculation is based on the interface columns that QoS uses to build scheduler profiles for L2TP sessions. For example, a typical interface column might consist of an L2TP session over an Ethernet VLAN over a Gigabit Ethernet interface.

You can configure QoS to control the rate of any logical interface in the interface column. For those logical interfaces with a rate controlled by QoS, QoS reports this configured rate as the transmit connect speed for that interface. For those logical interfaces that do not have a QoS-configured rate, QoS reports the speed of the underlying physical port as the transmit connect speed.

For more information, see *QoS and L2TP TX Speed AVP 24 Overview* in *JunosE Quality of Service Configuration Guide*.

### Actual

The actual method calculates the transmit connect speed of the subscriber's access interface as the lesser of the following two values:

- Value using the dynamic layer 2 calculation method
- Value using the QoS calculation method

## Transmit Connect Speed Calculation Examples

The examples in this section illustrate how the router uses the methods described in “[Transmit Connect Speed Calculation Methods](#)” on page 395 to calculate the transmit connect speed.

### Example 1: L2TP Session over ATM 1483 Interface

In this example, an L2TP session is established over an ATM 1483 subinterface on an OC3/STM1 ATM IOA. The configuration has the following characteristics:

- There is no explicit static configuration for the layer 2 (ATM 1483) interface.
- A transmit connect speed of 10 Mbps is provided dynamically from a RADIUS authentication server when the subscriber logs in.
- The transmit connect speed calculated by QoS is 5 Mbps.

Based on these characteristics, [Table 79 on page 397](#) lists the transmit connect speed value reported in L2TP Transmit (TX) Speed AVP 24 for each calculation method, and the reason why L2TP reports this value.

**Table 79: Transmit Connect Speeds for L2TP over ATM 1483 Example**

| Calculation Method | Transmit Connect Speed Reported in AVP 24 | Reason                                                                                                                                  |
|--------------------|-------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| Static layer 2     | 155 Mbps                                  | L2TP reports the speed of the underlying OC3 physical port because there is no explicit static configuration for the layer 2 interface. |
| Dynamic layer 2    | 10 Mbps                                   | L2TP reports the transmit connect speed provided by RADIUS.                                                                             |
| QoS                | 5 Mbps                                    | L2TP reports the transmit connect speed calculated by QoS.                                                                              |
| Actual             | 5 Mbps                                    | L2TP reports the lesser of the dynamic layer 2 speed (10 Mbps) or the QoS speed (5 Mbps).                                               |

### Example 2: L2TP Session over Ethernet VLAN Interface

In this example, an L2TP session is established over a PPPoE subinterface over an Ethernet VLAN subinterface. The configuration has the following characteristics:

- The Ethernet VLAN subinterface is configured with an advisory transmit speed of 100 Mbps.
- The dynamic layer 2 setting does not apply to the VLAN subinterface.
- The transmit connect speed calculated by QoS is 10 Mbps.

Based on these characteristics, [Table 80 on page 398](#) lists the transmit connect speed value reported in L2TP Transmit (TX) Speed AVP 24 for each calculation method, and the reason why L2TP reports this value.

**Table 80: Transmit Connect Speeds for L2TP over Ethernet Example**

| Calculation Method | Transmit Connect Speed Reported in AVP 24 | Reason                                                                                                                                                                                          |
|--------------------|-------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Static layer 2     | 100 Mbps                                  | L2TP reports the advisory transmit speed configured on the VLAN subinterface. If configured, the advisory transmit speed takes precedence over the physical port speed for a VLAN subinterface. |
| Dynamic layer 2    | 100 Mbps                                  | L2TP reports the static layer 2 value because the dynamic layer 2 setting does not apply to a VLAN subinterface.                                                                                |
| QoS                | 10 Mbps                                   | L2TP reports the transmit connect speed calculated by QoS.                                                                                                                                      |
| Actual             | 10 Mbps                                   | L2TP reports the lesser of the dynamic layer 2 speed (100 Mbps) or the QoS speed (10 Mbps).                                                                                                     |

## Transmit Connect Speed Reporting Considerations

The following considerations affect the transmit connect speed value reported in L2TP Transmit (TX) Speed AVP 24 when you use this feature.

### Session Termination for Dynamic Speed Timeout

Under certain heavy load conditions, the router might be unable to obtain the dynamic-layer2 value for the transmit connect speed of the subscriber's access interface. In this situation, the LAC sends the LNS an L2TP Call-Disconnect-Notify (CDN) message to terminate the L2TP session.

For more information about supported L2TP terminate reasons, see [“AAA Terminate Reasons” on page 265](#).

### Advisory Speed Precedence for VLANs over Bridged Ethernet

For interface columns that consist of an L2TP session over an Ethernet VLAN subinterface over a bridged Ethernet interface, the advisory transmit speed of the VLAN subinterface, if configured with the **vlan advisory-tx-speed** command, takes precedence over the physical port speed of the underlying layer 2 ATM 1483 interface. As a result, if the advisory

transmit speed is configured for the VLAN subinterface, L2TP reports this value as the transmit connect speed regardless of the port speed of the ATM 1483 interface.

## Using AAA Domain Maps to Configure the Transmit Connect Speed Calculation Method

To configure the transmit connect speed calculation method for a tunneled L2TP session associated with an AAA domain map:

1. Access Domain Map Tunnel Configuration mode.

```
host1(config)#aaa domain-map sunnyvale.com
host1(config-domain-map)#router-name lac
host1(config-domain-map)#tunnel 5
host1(config-domain-map-tunnel)#
```

For more information about how to map a domain to an L2TP tunnel from Domain Map Tunnel Configuration mode, see [“Mapping a User Domain Name to an L2TP Tunnel Overview” on page 347](#).

2. From Domain Map Tunnel Configuration mode, configure the calculation method for the transmit connect speed of the subscriber's access interface.

```
host1(config-domain-map-tunnel)#tx-connect-speed-method dynamic-layer2
```

3. (Optional) Use the **show aaa domain-map** command to verify configuration of the transmit connect speed calculation method.

```
host1(config-domain-map-tunnel)#run show aaa domain-map
```

```
Domain: sunnyvale.com; router-name: lac; ipv6-router-name: default
```

| Tunnel Tag | Tunnel Peer     | Tunnel Source         | Tunnel Type            | Tunnel Medium  | Tunnel Password | Tunnel Id             | Tunnel Client Name |
|------------|-----------------|-----------------------|------------------------|----------------|-----------------|-----------------------|--------------------|
| 5          | <null>          | <null>                | l2tp                   | ipv4           | <null>          | <null>                | <null>             |
| Tunnel Tag | Server Name     | Tunnel Preference     | Tunnel Max Sessions    | Tunnel RWS     |                 | Tunnel Virtual Router |                    |
| 5          | <null>          | 2000                  | 0                      | system chooses |                 | <null>                |                    |
| Tunnel Tag | Failover Resync | Tunnel Switch Profile | Tunnel Tx Speed Method |                |                 |                       |                    |
| 5          | <null>          | <null>                | dynamic layer2         |                |                 |                       |                    |

## Using AAA Tunnel Groups to Configure the Transmit Connect Speed Calculation Method

To configure the transmit connect speed calculation method for a tunneled L2TP session associated with an AAA tunnel group:

1. Access Tunnel Group Tunnel Configuration mode.

```
host1(config)#aaa tunnel-group boston
host1(config-tunnel-group)#tunnel 3
host1(config-tunnel-group-tunnel)#
```

For more information about how to map a domain to an L2TP tunnel from Tunnel Group Tunnel Configuration mode, see [“Mapping a User Domain Name to an L2TP Tunnel Overview” on page 347](#).

- From Tunnel Group Tunnel Configuration mode, configure the calculation method for the transmit connect speed of the subscriber's access interface.

```
host1(config-tunnel-group-tunnel)#tx-connect-speed-method qos
```

- (Optional) Use the **show aaa tunnel-group** command to verify configuration of the transmit connect speed calculation method.

```
host1(config-tunnel-group-tunnel)#run show aaa tunnel-group
```

```
Tunnel Group: boston
```

| Tunnel Tag | Tunnel Peer            | Tunnel Source         | Tunnel Type            | Tunnel Medium  | Tunnel Password | Tunnel Id             | Tunnel Client Name |
|------------|------------------------|-----------------------|------------------------|----------------|-----------------|-----------------------|--------------------|
| 3          | <null>                 | <null>                | l2tp                   | ipv4           | <null>          | <null>                | <null>             |
| Tunnel Tag | Tunnel Server Name     | Tunnel Preference     | Tunnel Max Sessions    | Tunnel RWS     |                 | Tunnel Virtual Router |                    |
| 3          | <null>                 | 2000                  | 0                      | system chooses |                 | <null>                |                    |
| Tunnel Tag | Tunnel Failover Resync | Tunnel Switch Profile | Tunnel Tx Speed Method |                |                 |                       |                    |
| 3          | <null>                 | <null>                | qos                    |                |                 |                       |                    |

## Using AAA Default Tunnel Parameters to Configure the Transmit Connect Speed Calculation Method

You can configure the transmit connect speed calculation method as a default AAA tunnel parameter by using the **aaa tunnel tx-connect-speed-method** command from Global Configuration mode. This command applies the specified calculation method to all tunneled L2TP sessions associated with a particular virtual router, and thereby alleviates the need for you to configure the transmit connect speed calculation method for each individual subscriber.

Configuring the calculation method as a default AAA tunnel parameter for a virtual router has lower precedence than using AAA domain maps, AAA tunnel groups, or RADIUS to configure the transmit connect speed calculation method. The router uses the calculation method specified with the **aaa tunnel tx-connect-speed-method** command if the tunnel attributes returned from an AAA domain map, an AAA tunnel group, or a RADIUS authentication server do not include the transmit connect speed calculation method.

To configure the transmit connect speed calculation method for all tunneled L2TP sessions associated with a particular virtual router:

- Create the virtual router for which you want to configure the transmit connect speed calculation method.

```
host1(config)#virtual-router north
```



For more information about configuring and using virtual routers, see the *Configuring Virtual Routers* chapter in the *JunosE System Basics Configuration Guide*.

2. Configure the transmit connect speed calculation method in the context of this virtual router.

```
host1:north(config)#aaa tunnel tx-connect-speed-method qos
```

- To specify the calculation method for the transmit connect speed, use one of the following keywords, as described in [“Using AAA Tunnel Groups to Configure the Transmit Connect Speed Calculation Method”](#) on page 399:

- **static-layer2**
- **dynamic-layer2**
- **qos**
- **actual**

3. (Optional) Use the **show aaa tunnel-parameters** command to verify configuration of the transmit connect speed calculation method.

```
host1:north(config)#run show aaa tunnel-parameters
Tunnel password is <NULL>
Tunnel client-name is <NULL>
Tunnel nas-port-method is none
Tunnel switch-profile is boston
Tunnel tx-connect-speed-method is qos
Tunnel nas-port ignore disabled
Tunnel nas-port-type ignore disabled
Tunnel assignmentId format is assignmentId
Tunnel calling number format is fixed
```

## Using RADIUS to Configure the Transmit Connect Speed Calculation Method

On the LAC, the router can receive tunnel configuration attributes through a RADIUS authentication server. To use RADIUS to configure the transmit connect speed calculation method for a subscriber's access interface, you can configure RADIUS to include the Tunnel-Tx-Speed-Method RADIUS attribute (Juniper Networks VSA 26-94) in RADIUS Access-Accept messages.

[Table 81 on page 402](#) describes the Tunnel-Tx-Speed-Method RADIUS attribute. For more information about RADIUS Access-Accept messages, see [“Subscriber AAA Access Messages Overview”](#) on page 182. For a description of the RADIUS attributes supported by JunosE Software, see [“RADIUS IETF Attributes”](#) on page 241.

Table 81: Tunnel--Tx-Speed-Method RADIUS Attribute

| Attribute Number | Attribute Name         | Description                                                                                                  | Length | Subtype Length | Value                                                                                                                                                                                                                                                                                                                                          |
|------------------|------------------------|--------------------------------------------------------------------------------------------------------------|--------|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| [26-94]          | Tunnel-Tx-Speed-Method | The method that the router uses to calculate the transmit connect speed of the subscriber's access interface | 12     | 6              | integer: <ul style="list-style-type: none"> <li>1 = static-layer2; TX speed based on static layer 2 settings</li> <li>2 = dynamic-layer2; TX speed based on dynamic layer 2 settings</li> <li>3 = qos; TX speed based on QoS settings</li> <li>4 = actual; TX speed that is the lesser of the dynamic-layer2 value or the qos value</li> </ul> |

#### Related Documentation

- [Transmit Connect Speed Calculation Methods on page 395](#)
- [Using AAA Domain Maps to Configure the Transmit Connect Speed Calculation Method on page 399](#)
- [Using AAA Tunnel Groups to Configure the Transmit Connect Speed Calculation Method on page 399](#)
- [Using AAA Default Tunnel Parameters to Configure the Transmit Connect Speed Calculation Method on page 400](#)
- [Using RADIUS to Configure the Transmit Connect Speed Calculation Method on page 401](#)
- *aaa tunnel tx-connect-speed-method*
- *tx-connect-speed-method*

## Controlling the Handling of L2TP Hello Packets in Fast Path

When the E series router acting as an LNS receives L2TP hello requests from the L2TP access concentrator (LAC), the LNS processes them by using an active switch route processor (SRP) module and returns L2TP hello responses to the LAC. When a unified ISSU operation is performed on the LNS deployed in an environment with more than 16,000 L2TP tunnels, the L2TP hello request may be timed out and lost for the ES2 10G UPLINK LM, ES2 10G LM, and ES2 10G ADV LM modules as these modules take more than 80 seconds to reload. This causes the L2TP tunnels established over ES2 10G UPLINK LM, ES2 10G LM, and ES2 10G ADV LM modules to be terminated.

To avoid such unexpected termination of L2TP tunnels, you can enable the handling of L2TP hello packets in fast path by using the *l2tp fastpath-hello-processing* command in the LNS. When the handling of L2TP hello packets in fast path is enabled, following steps are performed in the LNS during the unified ISSU operation:

1. Active SRP module informs the ES2 10G UPLINK LM, ES2 10G LM, and ES2 10G ADV LM modules acting as access line modules to divert L2TP hello requests to ES2 4G

LMs that act as Tunnel Service line modules (TSMs). The SRP module passes the details about all active TSMs to the access line modules.

2. Access line modules forward the hello requests received from the LAC to the TSMs in a round robin manner for load balancing.
3. TSMs process the received hello requests in their forwarding layer and forward the hello responses to the LAC directly from their forwarding layer.

Once the unified ISSU operation is completed, the active SRP module processes the L2TP hello requests.



**NOTE:**

- This feature is supported only on the E120 or E320 router with ES2 10G UPLINK LM, ES2 10G LM, and ES2 10G ADV LM modules acting as access line modules. This feature is not supported for ES2 10G ADV LM acting as TSM.
- If an ES2 4G LM acting as TSM goes down during unified ISSU, the round robin mechanism is not modified dynamically.
- The L2TP hello request is processed only by the ES2 4G LM acting as shared tunnel server or the ES2 4G LM with ES2-S1 SERVICE IOA.

To enable the handling of L2TP hello packets in fast path:

- Issue the **l2tp fastpath-hello-processing** command in Global Configuration mode  
`host1(config#) l2tp fastpath-hello-processing`

Use the **no** version to disable the handling of L2TP hello packets in fast path.



**CAUTION:** This feature may cause the L2TP tunnel hello packet to be dropped during peak tunneled data traffic conditions.

**Related Documentation**

- [Monitoring Global Configuration Status on E Series Routers on page 433](#)
- [Unexpected Termination of L2TP Tunnels after Unified ISSU](#)

## PPP Accounting Statistics

JunosE accounting for tunneled subscribers at the L2TP LAC counts the payload that PPP passes to or receives from L2TP for transport. At this stage in the protocol processing, any padding outside PPP, such as that for PPPoE, has been removed. Accounting includes the authentication acknowledgement packet, CHAP success packets, and PAP

acknowledgment packets. Accounting ends when L2TP has been notified to terminate the session. The statistics are reported in the following RADIUS attributes:

| Attribute Number | Attribute Name      |
|------------------|---------------------|
| 42               | Acct-Input-Octets   |
| 43               | Acct-Output-Octets  |
| 47               | Acct-Input-Packets  |
| 48               | Acct-Output-Packets |

Termination of a tunneled session can result from PPP termination, L2TP shutdown, subscriber logout, or lower layer down events. When the session is terminated through PPP, the software counts both the PPP terminate-request and the PPP terminate-acknowledgement packets.

- Accounting statistics reported in RADIUS octet counts (Acct-Input-Octets and Acct-Output-Octets) for tunneled PPP customers at the L2TP LAC include the following data:
  - All upper layer control traffic, including IPCP, IPCPv6, OSICP, and MPLSNCP
  - All data traffic, including IP, IPv6, MPLS, and OSI
  - PPP PAP or CHAP acknowledgments, and also retransmission of PAP or CHAP that take place after the session is active (even when proxy authentication is accepted)
  - All PPP PAP or CHAP negotiations in the case where proxy authentication is disabled or required to renegotiate at the LNS
  - All LCP traffic when proxy LCP is disabled or required to renegotiate at the LNS
  - All PPP LCP echo requests and their responses
  - PPP LCP terminate-request or terminate-acknowledgement packets from the client or LNS when PPP initiates termination of the session
  - If present, the two PPP header bytes (Address Field 0xFF and Control Field 0x03) as part of the L2TP payload
- Accounting statistics reported in RADIUS octet counts (Acct-Input-Octets and Acct-Output-Octets) for tunneled PPP customers at the L2TP LAC exclude the following data:
  - LCP when Proxy LCP is enabled and accepted at the LNS
  - Initial PPP PAP request
  - Initial PPP CHAP challenge and response
- Accounting statistics reported in RADIUS packet counts (Acct-Input-Packets and Acct-Output-Packets) for tunneled PPP customers at the L2TP LAC are based on

packets delivered to or received from the L2TP session. These statistics exclude L2TP control traffic and L2TP hello messages.

For information on accounting statistics for terminated PPP sessions, see the *PPP Accounting Statistics Overview* section in the *JunosE Link Layer Configuration Guide*.

#### Related Documentation

- [Application Support for Stateful Line Module Switchover](#)
- [Collecting Accounting Statistics on page 69](#)
- [RADIUS IETF Attributes Supported for Subscriber AAA Accounting Messages on page 193](#)

## Stateful Line Module Switchover for LNS Sessions

In releases in which the stateful line module switchover feature is not available or in scenarios in which this behavior is disabled, a reload of the line module disconnects user sessions and disrupts traffic forwarding through it. In a network in which an E120 or E320 router that contains the Service IOA functions as the LNS device on one side of the L2TP tunnel, the LNS is the logical termination point of a PPP connection that is being tunneled from the remote system by the LAC. A LAC receives packets from a remote client and forwards them to an LNS on a remote network. All the tunneled sessions terminate on the LNS to provide enhanced performance during decapsulation and encapsulation of packets, and fragmentation and reassembly of tunneled packets. If the line module in the LNS that performs the traffic processing encounters a fault, such as a hardware or software error, all the active subscriber sessions are disconnected.

Stateful switchover of LNS sessions avoids subscriber disconnections during the switchover of the line module installed on the LNS device (tunnel server module or ES2-S1 Service IOA on ES2 4G LMs in this case). You can enable high availability for the line module pairs using the **mode high-availability slot** command in Redundancy Configuration mode. This command enables you to specify the slots in which the tunnel server line modules that you want to be configured as the primary and secondary modules reside. If HA is active between these modules, the secondary module becomes the primary when the assigned primary module fails. The newly active primary module retains all the subscribers that were active and were managed by the previously configured primary module without requiring the subscribers to be reconnected. The failure of the tunnel server module in the LNS device and the switchover from a defective module to a newly active primary module in a seamless, uninterrupted manner for subscribers is transparent to the end users.

Line module high availability uses a 1:1 redundancy model to maintain subscriber sessions, and this functionality is supported only on E120 and E320 routers installed with ES2 4G LMs and Service IOAs. This feature is supported only for PPP-based stacks (such as L2TP, PPP, and IP) and not for other applications such as GRE.

The router uses the tunnel server module to increase the performance of packet processing by offloading the decapsulation and reassembly of packets to the tunnel server module. All the L2TP and PPP session data are downloaded to the tunnel server module to assist this operation. When the primary tunnel server module fails, either due to hardware or software error, subscribers are disconnected because of the PPP keepalive

expiry mechanism and also because the forwarding path is not maintained. When stateful switchover for LNS sessions is enabled, you can provision another tunnel server module as the secondary module in 1:1 mode. When this feature is enabled, all the required session data is mirrored to the secondary module. Any session data change, such as session creation or deletion, is mirrored from the primary to the secondary module. The previously configured primary module, after it becomes operational, takes over the role of the secondary module.

**Related  
Documentation**

- *Stateful Line Module Switchover Overview*
- *Preservation of Statistics During Stateful Line Module Switchover*
- *Application Support for Stateful Line Module Switchover*

## CHAPTER 15

# Configuring L2TP Dial-Out

This chapter describes the Layer 2 Tunneling Protocol (L2TP) dial-out feature on your E Series router. This chapter includes the following sections:

- [L2TP Dial-Out Overview on page 407](#)
- [L2TP Dial-Out Terms on page 408](#)
- [L2TP Dial-Out Platform Considerations on page 409](#)
- [L2TP Dial-Out References on page 409](#)
- [L2TP Dial-Out Network Model on page 409](#)
- [L2TP Dial-Out Process on page 410](#)
- [L2TP Dial-Out Operational States on page 411](#)
- [L2TP Dial-Out Outgoing Call Setup Details on page 414](#)
- [Creating a Profile Before Configuring L2TP Dial-Out on page 416](#)
- [Configuring L2TP Dial-Out on page 417](#)

### L2TP Dial-Out Overview


---

L2TP dial-out provides a way for corporate virtual private networks (VPNs) that use Broadband Remote Access Server (B-RAS) to dial out to remote offices that have only narrowband dial-up access. The L2TP network server (LNS) function is deployed in networks that have a combination of broadband and narrowband access.

A remote site can communicate on demand with the home site with a normal L2TP access concentrator (LAC) to LNS session. When the communication finishes, the remote site terminates the session. However, if the home site wishes to communicate with the remote site and no incoming call is currently established, the home site needs a method to dial out to the remote site. This method is L2TP dial-out, which uses the L2TP outgoing call support defined in RFC 2661—Layer Two Tunneling Protocol “L2TP” (August 1999).

[Figure 10 on page 408](#) shows the dial-out model in which the LNS initiates L2TP sessions and provides enough information to the narrowband LAC so that it can complete the dial-out from the home site to the remote site.

Figure 10: Network Model for Dial-Out



**NOTE:** The dial-out feature exists in the LNS only. It does not exist in the LAC.

#### Related Documentation

- [L2TP Overview on page 321](#)
- [L2TP Dial-Out Terms on page 408](#)
- [L2TP Dial-Out Network Model on page 409](#)
- [L2TP Dial-Out Operational States on page 411](#)
- [L2TP Dial-Out Process on page 410](#)
- [L2TP Dial-Out Outgoing Call Setup Details on page 414](#)

## L2TP Dial-Out Terms

Table 82 on page 408 describes key terms used in L2TP dial-out.

**Table 82: L2TP Dial-Out Terms**

| Term             | Description                                                                                                                                                                                                                                                                                          |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Dial-out trigger | IP packet that initiates a dial-out session                                                                                                                                                                                                                                                          |
| Dial-out session | Control entity for a triggered IP flow used to manage the establishment of an associated L2TP session for dial-out                                                                                                                                                                                   |
| Dial-out target  | A virtual router context and an IP address prefix, for which the arrival of an IP packet (a dial-out trigger) initiates a dial-out session.                                                                                                                                                          |
| Dial-out route   | Contains the dial-out target, as well as a domain name and profile. <ul style="list-style-type: none"> <li>• The domain name is used in the initial Access-Request message.</li> <li>• The profile is used to create the IP/Point-to-Point Protocol (PPP) stack for the dial-out session.</li> </ul> |

#### Related Documentation

- [L2TP Dial-Out Overview on page 407](#)
- [L2TP Dial-Out Network Model on page 409](#)
- [L2TP Dial-Out Operational States on page 411](#)



- [L2TP Dial-Out Process on page 410](#)
- [L2TP Dial-Out Outgoing Call Setup Details on page 414](#)

## L2TP Dial-Out Platform Considerations

---

L2TP dial-out is supported on all E Series routers.

For information about the modules supported on E Series routers:

- See the *ERX Module Guide* for modules supported on ERX7xx models, ERX14xx models, and the ERX310 Broadband Services Router.
- See the *E120 and E320 Module Guide* for modules supported on the E120 and E320 Broadband Services Routers.

### Related Documentation

- [L2TP Dial-Out Overview on page 407](#)
- [L2TP Dial-Out Network Model on page 409](#)

## L2TP Dial-Out References

---

For more information about L2TP, see RFC 2661—Layer Two Tunneling Protocol “L2TP” (August 1999).

### Related Documentation

- [L2TP Dial-Out Overview on page 407](#)
- [L2TP Dial-Out Network Model on page 409](#)

## L2TP Dial-Out Network Model

---

In the figure in “[L2TP Dial-Out Overview](#)” on [page 407](#), the home site connects to the Internet over a permanent leased line to the Internet service provider’s (ISP’s) E Series LNS. The ISP uses an IP network to connect the LNS to the narrowband access point of the network where the narrowband LAC exists. The narrowband LAC connects to a narrowband network (ISDN) that the remote site is also connected to.

The figure shows three RADIUS servers. The home site maintains the home server, and the other two servers are at the LNS and the LAC. The router accesses the home and LNS RADIUS servers. (The separation of the RADIUS servers is transparent to the router.)

Before any attempts at connectivity can take place from the home site to the remote site, an administrator must configure a dial-out route on the router. This route directs the router to start a dial-out operation. The route includes a dial-out target (the virtual router context and the IP address of the remote site). When the router receives a packet destined for the target, it triggers a dial-out session to the target. The route is associated with a profile that holds parameters for the interface stack that the router builds as a result of the dial-out.

- Related Documentation**
- [L2TP Dial-Out Overview on page 407](#)
  - [L2TP Dial-Out Terms on page 408](#)
  - [L2TP Dial-Out Operational States on page 411](#)
  - [L2TP Dial-Out Process on page 410](#)
  - [L2TP Dial-Out Outgoing Call Setup Details on page 414](#)

---

## L2TP Dial-Out Process

The following is the dial-out process used in the network model illustrated in “[L2TP Dial-Out Overview](#)” on page 407:

1. The router receives a trigger packet.
2. The router builds a RADIUS Access-Request message and sends it to the RADIUS server that is associated with the virtual router on which the dial-out route is defined—typically, the RADIUS home server.
3. The RADIUS server’s response to the Access-Request is similar to the response used for LAC incoming calls. Notable differences are that the IP addresses of the peer are interpreted as LAC addresses instead of LNS addresses. In addition, narrowband details, such as calling numbers, are returned.
4. The LNS makes the outgoing call using a load-balancing or round-robin mechanism identical to the one that the E Series LAC uses for incoming calls. The LAC may also employ the LAC RADIUS in tunnel authentication.
5. Once the LNS successfully completes a control connection and session with the LAC, the LAC performs the actual narrowband dial-out operation to the remote site using the information passed by the LNS during session setup.
6. A PPP session is started on the remote customer premises equipment (CPE), and mutual PPP authentication is performed at the remote CPE and the LNS as follows:
  - a. The LNS uses the LNS RADIUS server to validate the remote CPE’s PPP session, while the CPE can use its own RADIUS server to validate the LNS’s PPP session.
  - b. The LNS uses the username and password that is returned in the first Access-Accept message.
7. Once authentication is successful, an IP interface is built on top of the PPP interface at the LNS. Internet Protocol Control Protocol (IPCP) is negotiated, and the framed route that RADIUS returns as a result of the PPP authentication supersedes the dial-out route.

IP traffic can now flow freely between the home and remote sites.

- Related Documentation**
- [L2TP Dial-Out Overview on page 407](#)
  - [L2TP Dial-Out Terms on page 408](#)
  - [L2TP Dial-Out Network Model on page 409](#)

- [L2TP Dial-Out Operational States on page 411](#)
- [L2TP Dial-Out Outgoing Call Setup Details on page 414](#)

## L2TP Dial-Out Operational States

The dial-out state machine is a control process within the router that manages the dial-out function for each IP flow. The dial-out state machine has four levels of control: the router chassis, virtual router, targets, and sessions. This section describes the operational states of each of these levels.

### Chassis

[Table 83 on page 411](#) describes the operational states of the chassis.

**Table 83: Chassis Operational States**

| State                | Description                                                                                                                                                 |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| inService            | Dial-out service is operational at the chassis level.                                                                                                       |
| initializationFailed | Dial-out service could not obtain enough system resources for basic operation. All configuration commands fail, and the dial-out service does not function. |

### Virtual Router

[Table 84 on page 411](#) describes the operational states of the virtual router.

**Table 84: Virtual Router Operational States**

| State       | Description                                                                                                                         |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------|
| inService   | Dial-out service is operational for the virtual router.                                                                             |
| initPending | Dial-out service is waiting for the virtual router to be operational. Targets defined within the virtual router are not functional. |
| down        | The dial-out interface for this virtual router is down. Targets defined within the virtual router are not functional.               |

### Targets

[Table 85 on page 411](#) describes the operational states of the targets.

**Table 85: Target Operational States**

| State     | Description                           |
|-----------|---------------------------------------|
| inService | Dial-out route is up and operational. |

**Table 85: Target Operational States** (*continued*)

| State     | Description                                                                                                                                                                                                                                                                                                                                                               |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| inhibited | <p>Dial-out service cannot obtain sufficient resources to handle triggers, and all triggers are discarded. When resources become available, a target can transition from inhibited to inService.</p> <p>Note that sessions within an inhibited target that are already in the process of connecting or are in the inService state are not affected by this condition.</p> |
| down      | <p>There are insufficient resources to support the creation of a dial-out route for the target. When resources become available, the target can transition to inService.</p> <p>Note that sessions within a down target that are already in the process of connecting or are in the inService state are not affected by this condition.</p>                               |

## Sessions

Table 86 on page 412 describes operational states of the sessions.

**Table 86: Session Operational States**

| State          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| authenticating | <p>New sessions start in the authenticating state. In this state, the dial-out state machine has received a valid trigger and is waiting for authentication, authorization, and accounting (AAA) to complete the initial authentication.</p> <p>On getting a grant from AAA, the session transitions to the connecting state. Alternatively, on getting a deny from AAA, the session transitions to the inhibited state.</p>                                                      |
| connecting     | <p>Sessions enter the connecting state when authentication is complete. In this state, the dial-out state machine has initiated an outgoing L2TP call. On entering this state, the session-connecting timer is set to the chassis-wide trigger timer value. The session stays in this state until either the outgoing call is successful or the connecting timer expires. Any new trigger packets received for this session when it is in the connecting state are discarded.</p> |
| inService      | <p>A session enters the inService state from the connecting state on successful completion of the dial-out call request. The session stays in this state until the outgoing call is closed.</p>                                                                                                                                                                                                                                                                                   |

Table 86: Session Operational States (*continued*)

| State         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| inhibited     | <p>A session enters the inhibited state from the connecting state when the connecting timer expires (that is, the outgoing call was unsuccessful). This state prevents the router from thrashing on an outgoing call that cannot be completed. When in this state, the router discards all trigger packets received for the session.</p> <p>The inhibited timer controls the amount of time spent in this state. The setting of the inhibited timer varies depending on whether the session is entering the inhibited state for the first time or is reentering the state.</p> <ul style="list-style-type: none"> <li>• If it is the first time, the inhibited timer is initialized to the chassis-wide trigger value.</li> <li>• If it is reentering the state, the inhibited timer is initialized to 2 times the previous value of the inhibited timer, up to a maximum of 8 times the chassis-wide trigger value. For example, if the chassis-wide trigger value is 30 seconds, the setting of the inhibited timer within the session (on subsequent immediate reentries; see postInhibited state) is 30, 60, 120, 240. Since 240 is 8 x 30, the inhibited timer for this session is never set larger than 240 seconds.</li> </ul> |
| postInhibited | <p>A session enters the postInhibited state after completion of an inhibited state. The inhibited timer is reused to control the amount of time the session stays in postInhibited state. In this state the timer repeatedly times out and reduces the inhibited timer by a factor of 2 on each iteration. Once the inhibited timer reaches zero, the session transitions to dormant. The receipt of a trigger in this state results in a transition to the authenticating state.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| dormant       | <p>A session enters the dormant state after completion of a postInhibited state. The dormant timer is initialized to the chassis-wide dormant timer value, minus the time the session spent in the postInhibited state. Receipt of a new trigger packet transitions the session to the authenticating state. If the dormant timer expires, the session is deleted. The dormant state exists to allow analysis of a dial-out session before it is deleted.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| pending       | <p>A session enters the pending state when a valid trigger is received but there already are the maximum number of connecting sessions in the router. The router discards all subsequent trigger packets until other sessions transition out of the connecting state. When this happens, pending sessions can transition to the dormant state.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| failed        | <p>A session enters the failed state when the router detects a configuration error that prevents the successful operation of the session. Specifically, one of the final steps in a dial-out request is mutual PPP authentication at the LNS. A side-effect of authentication is the installation of an access route for the outgoing call. If the access route does not correspond to the trigger packet (that is, the trigger packet cannot be routed successfully by the new access route), the router detects this discrepancy as a configuration error because trigger packets that arrive are not forwarded into the outgoing call; rather, they are buffered or discarded.</p> <p>The only way to exit the failed state is with the <b>l2tp dial-out session reset</b> command.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                            |

- Related Documentation**
- [L2TP Dial-Out Overview on page 407](#)
  - [L2TP Dial-Out Terms on page 408](#)
  - [L2TP Dial-Out Network Model on page 409](#)
  - [L2TP Dial-Out Process on page 410](#)
  - [L2TP Dial-Out Outgoing Call Setup Details on page 414](#)

## L2TP Dial-Out Outgoing Call Setup Details

This section details the process described in “[L2TP Dial-Out Process](#)” on page 410.

### Access-Request Message

To create the username in the authentication request, the router uses the trigger, dial-out route, domain name, and optional Multiprotocol Label Switching (MPLS) route distinguisher (RD). The username is constructed as follows:

***[MPLS RD]/{trigger destination address}@domain-name***

For example, given a dial-out route with an IP prefix of 10.10.0.0/16, a domain name of L2TP-dial-out.de.dt, and an MPLS RD of 0.0.0.0:65000, if a trigger packet arrives with a destination IP address of 10.10.1.1, the router creates the following username:

**0.0.0.0:65000/10.10.1.1@L2TP-dial-out.de.dt**

No password is offered, and the authentication request is passed to the S-series AAA server for normal authentication processing.

Using the above example, the AAA domain map processes the L2TP-dial-out.de.dt domain as for any other domain. If RADIUS authentication is configured for the authenticating virtual router (VR) context, AAA passes the authentication request to the E Series RADIUS client. The RADIUS authentication request is consistent with other requests, except that the Service-Type attribute is set to outbound (value of 5).

### Access-Accept Message

The router expects RADIUS attributes that define a tunnel to be returned with the additions in [Table 87 on page 414](#). If tunnel attributes are excluded from the Access-Accept message or the returned Service-Type attribute is not set to outbound, the dial-out session is denied.

**Table 87: Additions to RADIUS Attributes in Access-Accept Messages**

| Attribute Number  | Attribute Name         | Content              |
|-------------------|------------------------|----------------------|
| 6                 | Service-Type           | Outbound             |
| 67                | Tunnel-Server-Endpoint | IP address of LAC    |
| Juniper VSA 26-35 | Tunnel-Dialout-Number  | L2TP dial-out number |

**Table 87: Additions to RADIUS Attributes in Access-Accept Messages (continued)**

| Attribute Number  | Attribute Name     | Content                                                                                                                                |
|-------------------|--------------------|----------------------------------------------------------------------------------------------------------------------------------------|
| Juniper VSA 26-36 | PPP-Username       | Username used in PPP L2TP dial-out sessions at the LNS                                                                                 |
| Juniper VSA 26-37 | PPP-Password       | Password used in PPP L2TP dial-out sessions at the LNS                                                                                 |
| Juniper VSA 26-38 | PPP-Protocol       | Authentication protocol used for L2TP sessions.<br><br>0 = none<br><br>1 = PAP<br><br>2 = CHAP<br><br>3 = PAP-CHAP<br><br>4 = CHAP-PAP |
| Juniper VSA 26-39 | Tunnel-Min-Bps     | Minimum line speed; passed to LAC (not interpreted by the LNS)                                                                         |
| Juniper VSA 26-40 | Tunnel-Max-Bps     | Maximum line speed; passed to LAC (not interpreted by the LNS)                                                                         |
| Juniper VSA 26-41 | Tunnel-Bearer-Type | Bearer capability required: 0=name; 1=analog; 2=digital. Passed to LAC (not interpreted by the LNS).                                   |

## Outgoing Call

After receiving a valid tunnel definition from AAA, the E Series LNS initiates an outgoing call. The router follows the same load-sharing mechanisms as for incoming calls. See [“Configuring LAC Tunnel Selection Parameters” on page 361](#).

After an outgoing call is successfully signaled, the router dynamically creates a PPP interface. The profile in the dial-out route definition specifies any PPP configuration options. Both the L2TP session and the PPP interface exist on a Service module, identical to the LNS operation for incoming calls.

Once the PPP interface is created, Link Control Protocol (LCP) and IPCP are negotiated.

## Mutual Authentication

Mutual authentication takes place in LCP, where the LNS validates the PPP interface on the remote CPE and vice-versa. LNS takes the same actions to authenticate the peer as it does on incoming calls.

The LNS obtains the PPP username and password from the initial Access-Accept message. It then provides this information to the remote CPE for authentication.

## Route Installation

Once authentication is complete, the router creates a new access route. This route directs the forwarding of IP packets related to the original trigger packet to the newly created interface. The route does not need to be identical to the one specified in the dial-out route, but it must be able to forward packets that have the same destination address as the trigger packet. However, if the access route does not encompass the dial-out route definition, any other trigger packets initiate a new dial-out session.

The dial-out state machine verifies that the trigger packet can be forwarded over the route.

- If the verification is unsuccessful, the dial-out session is put into the failed state.
- If the verification is successful, the dial-out session is put into the inService state.

### Related Documentation

- [L2TP Dial-Out Overview on page 407](#)
- [L2TP Dial-Out Terms on page 408](#)
- [L2TP Dial-Out Network Model on page 409](#)
- [L2TP Dial-Out Operational States on page 411](#)
- [L2TP Dial-Out Process on page 410](#)

---

## Creating a Profile Before Configuring L2TP Dial-Out

Create a profile that the router uses to create the dynamic PPP and IP interfaces on the LNS. The profile specifies parameters that are common to all dial-out sessions that use the profile. The following is an example of a typical profile configuration.

1. Create a profile.

```
host1(config)#profile dialOut
host1(config-profile)#
```

2. Specify the interface used for dialout.

```
host1(config-profile)#ip unnumbered loopback 0/0
```

3. Specify the virtual router for the dial-out user's IP interface.

```
host1(config-profile)#ip virtual-router lns
```

4. Specify the authentication mechanism.

```
host1(config-profile)#ppp authentication chap
```

### Related Documentation

- [L2TP Dial-Out Network Model on page 409](#)
- [L2TP Dial-Out Process on page 410](#)
- [L2TP Dial-Out Outgoing Call Setup Details on page 414](#)
- *ip virtual-router*



- *ip unnumbered*
- *ppp authentication*
- *profile*

## Configuring L2TP Dial-Out

L2TP dial-out configuration tasks include the following sets of tasks:

- [Creating an L2TP Dial-Out Session on page 417](#)
- [Specifying the Maximum Timeout Period for Establishing an L2TP Dial-Out Session on page 417](#)
- [Specifying the Duration for an L2TP Dial-Out Session to Remain in Dormant State on page 418](#)
- [Specifying the Maximum Triggers to Buffer for an L2TP Dial-Out Session on page 418](#)
- [Deleting an L2TP Dial-Out Session on page 418](#)
- [Resetting an L2TP Dial-Out Session on page 419](#)

### Creating an L2TP Dial-Out Session

You can define an L2TP dial-out target by using the **l2tp dial-out target** command. When the router receives packets destined for the target, it creates a dial-out session.

To create a dial-out session:

- Issue the **l2tp dial-out target** command in Global Configuration mode.
- ```
host1(config)#l2tp dial-out target 10.10.0.0 255.255.0.0 L2TP-dial-out.de.dt profile dialOut
```

When you create a target, you must specify the following:

- *ipAddress*—IP address of the target
- *ipAddressMask*—IP address mask of the target
- *domainName*—Domain name used in the outgoing call Access-Request message
- *profileName*—Name of profile used to create the interface stack

Use the **default** version to remove the L2TP dial-out route. Use the **no** version to remove the L2TP dial-out route or target.

Specifying the Maximum Timeout Period for Establishing an L2TP Dial-Out Session

You can optionally set the maximum time allowed for successful establishment of an L2TP dial-out session.

To set the maximum time allowed for attempts to establish L2TP dial-out sessions:

- Issue the **l2tp dial-out connecting-timer-value** command in Global Configuration mode.

host1(config)#l2tp dial-out connecting-timer-value 30

If the session fails to be established before the connecting timer expires, subsequent attempts to establish the dial-out session to the same destination are inhibited temporarily. The range is 30–3600 seconds.

Use the **no** version to set the connecting timer to the default, 30 seconds.

Specifying the Duration for an L2TP Dial-Out Session to Remain in Dormant State

You can optionally specify the duration for which the dial-out session stays in the dormant state waiting for a new trigger after the associated L2TP outgoing call ends.

To set how long the dial-out session waits in the dormant state for a new trigger after the associated L2TP outgoing call ends:

- Issue the **l2tp dial-out dormant-timer-value** command in Global Configuration mode.

host1(config)#l2tp dial-out dormant-timer-value 300

If no trigger is received before the dormant timer expires, the dial-out session is deleted. The range is 0–3600 seconds.

Use the **no** version to set the dormant timer to the default, 300 seconds (5 minutes).

Specifying the Maximum Triggers to Buffer for an L2TP Dial-Out Session

You can optionally set the maximum number of buffered trigger packets held for any dial-out session pending the successful establishment of the L2TP session. Once the session is established, the buffered trigger packets are transmitted.

To set the maximum number of trigger packets held in buffer while the dial-out session is being established:

- Issue the **l2tp dial-out max-buffered-triggers** command in Global Configuration mode.

host1(config)#l2tp dial-out max-buffered-triggers 50

If the configured maximum number of buffered trigger packets is exceeded, any new trigger packets received are discarded. The range of values is 0–50.

Use the **no** version to set the number of trigger buffers to the default, 0.

Deleting an L2TP Dial-Out Session

You can manually delete a dial-out session to close any L2TP outgoing call associated with the dial-out session.

To delete a dial-out session:

- Issue the **l2tp dial-out session delete** command in Global Configuration mode.

host1#l2tp dial-out session delete 10.10.0.0

There is no **no** version.

Resetting an L2TP Dial-Out Session

You can reset a dial-out session by forcing it to the dormant state. After you reset a dial-out session, any L2TP outgoing call associated with that session is closed.

To force the dial-out session to the dormant state where it remains until the dormant timer expires or it receives a new trigger:

- Issue the **l2tp dial-out session reset** command in Global Configuration mode.

```
host1#l2tp dial-out session reset 10.10.0.0
```

There is no **no** version.

Related Documentation

- [L2TP Dial-Out Network Model on page 409](#)
- [L2TP Dial-Out Process on page 410](#)
- [L2TP Dial-Out Outgoing Call Setup Details on page 414](#)
- [Monitoring Chassis-wide Configuration for L2TP Dial-out on page 449](#)
- [Monitoring Status of Dial-out Sessions on page 454](#)
- [Monitoring Dial-out Targets within the Current VR Context on page 455](#)
- [Monitoring Operational Status within the Current VR Context on page 456](#)
- *l2tp dial-out connecting-timer-value*
- *l2tp dial-out dormant-timer-value*
- *l2tp dial-out max-buffered-triggers*
- *l2tp dial-out session delete*
- *l2tp dial-out session reset*
- *l2tp dial-out target*

L2TP Disconnect Cause Codes

- [L2TP Disconnect Cause Codes on page 421](#)

L2TP Disconnect Cause Codes

Table 88 on page 421 describes the Point-to-Point Protocol (PPP) disconnect cause codes that are displayed by the **show l2tp received-disconnect-cause-summary** command, sorted by code number. For additional information, see RFC 3145.

Table 88: PPP Disconnect Cause Codes

Code	Name	Description
0	no info	<p>Code 0 includes disconnect causes that are not specifically identified by other codes. This code is generated in the following circumstances:</p> <ul style="list-style-type: none"> • Internal resource constraints (for example, excessive load or reduced resource availability) have prevented the generation of a more specific disconnect code. • RFC 3145 does not define a disconnect code that corresponds to the cause of the disconnection. <p>The following list shows current disconnection causes on an E Series LNS that do not have a specific disconnect cause codes:</p> <ul style="list-style-type: none"> • The peer initiated termination of LCP after the completion of LCP negotiations, but prior to proceeding to authentication of NCP negotiation. No conditions occurred that enabled the LNS to infer a more informative disconnect code. • The peer initiated renegotiation of LCP. • Invalid local MRU (for example, MRU negotiation has been disabled, but the lower MRU is less than the default MRU of 1500). • Unexpected local MLPPP MRRU for existing bundle (RFC 3145 code 10 covers peer MRRU mismatches, but not local mismatches). • Authentication failures not covered by any of the authentication-related codes (codes 13-16), such as: <ul style="list-style-type: none"> • Authentication denial of the local LCP by the peer • Local authentication failure due to no resources • Local authentication failure due to no authenticator

Table 88: PPP Disconnect Cause Codes (*continued*)

Code	Name	Description
1	admin disconnect	<p>The disconnection was a result of direct administrative action, including:</p> <ul style="list-style-type: none"> • The administrator shut down the network or link interface. • The administrator logged out the subscriber.
2	renegotiation disabled	Code 2 is not used; the E Series LNS is always capable of renegotiating LCP if proxy data is not available.
3	normal disconnect	<p>Indicates that one of the following events occurred:</p> <ul style="list-style-type: none"> • user-initiated logout (direction 1) • session timeout (direction 2) • inactivity timeout (direction 2) • address lease expired (direction 2) <p>The E Series LNS determines by inference that a normal disconnect has occurred for direction 1. The LNS does this when the peer initiates LCP termination after proceeding beyond the successful negotiation of LCP (that is, after starting authentication signaling or NCP negotiation).</p> <p>NOTE: The Error-code field is included by default in the Result Error Code attribute value pair (AVP) in L2TP Call-Disconnect-Notify (CDN) messages, even in normal disconnect cases when the peer initiates LCP termination after proceeding beyond LCP negotiation.</p>
4	compulsory encryption refused	<p>Code 4 with direction 2 is generated if the following conditions are met:</p> <ul style="list-style-type: none"> • The peer initiates LCP termination without having proceeded beyond the completion of LCP negotiation, and • Prior to receiving the terminate request from the peer, the local LCP has sent a Protocol Reject in response to any packet for Encryption Control Protocol (ECP) protocols (protocol codes 0x8053, 0x8055) from the peer. <p>Code 4 with direction 1 is never generated, because the E Series LNS never requests ECP.</p>
5	lcp failed to converge	An LCP configuration error prevented LCP from converging; the two peers attempted to negotiate but did not agree on acceptable LCP parameters.
6	lcp peer silent	LCP negotiation timed out; the LNS did not receive any LCP packets from the LAC.
7	lcp magic number error	A magic number error was detected; this indicates a possible looped back link.
8	lcp keepalive error	The keepalive drop count was exceeded.

Table 88: PPP Disconnect Cause Codes (*continued*)

Code	Name	Description
9	lcp mlppp endpoint discriminator mismatch	Code 9 is not used. Dynamic MLPPP bundling, which is the only kind of MLPPP bundling supported for MLPPP/L2TP, uses the endpoint discriminator as part of the key for bundle selection. Therefore, there will never be an unexpected endpoint discriminator for an existing MLPPP bundle.
10	lcp mlppp mrru not valid	The link attempted to join an existing MLPPP bundle whose peer maximum received reconstructed unit (MRRU) did not match the peer MRRU negotiated by the link.
11	lcp mlppp peer ssn invalid	Code 11 is not used; the short sequence number (SSN) option is not supported.
12	lcp callback refused	<p>Code 12 with direction 2 is generated when the following conditions are met:</p> <ul style="list-style-type: none"> • The peer initiates LCP termination without having proceeded to NCP negotiation, and • Prior to the termination, the local LCP has responded with a negative acknowledgement (NAK) to a callback option (LCP option 13) from the peer. <p>The E Series LNS never generates code 12 with direction 1 because the LNS never requests callback.</p>
13	authenticate timed out	Authentication failed because the authentication protocol timed out; either the CHAP Authenticate Response or the PAP Authenticate Request was not received.
14	authenticate mlppp name mismatch	Code 14 is not used. Dynamic MLPPP bundling, which is the only kind of MLPPP bundling supported for MLPPP/L2TP, uses the authenticated name as part of the key for bundle selection. Therefore, there will never be an unexpected authenticated name for an existing MLPPP bundle.
15	authenticate protocol refused	<p>No acceptable authentication protocol was negotiated by LCP.</p> <ul style="list-style-type: none"> • Code 15 with direction 1 is generated if the peer rejected all of the authentication protocols requested by the local LCP. • Code 15 with direction 2 is generated if the following conditions are met: <ul style="list-style-type: none"> • The peer initiates LCP termination without having proceeded beyond completion of NCP negotiation, and • During LCP negotiation, the local LCP responded with a NAK to the final authentication protocol requested by the peer.

Table 88: PPP Disconnect Cause Codes (*continued*)

Code	Name	Description
16	authenticate failure	<ul style="list-style-type: none"> Code 16 with direction 1 is generated if the local authentication of the peer fails (that is, the authenticator sent a PAP NAK or CHAP Failure packet) Code 16 with direction 2 is generated if the peer authentication of the local LCP fails (that is, the authenticator received a PAP NAK or CHAP Failure packet). <p>Note that there are a variety of causes for authentication failures, including bad credentials (bad name, password or secret) and resource problems.</p>
17	ncp no negotiation completed	<p>Code 17 is generated only if an NCP configuration error has prevented NCP negotiation from converging. This occurs when the two peers do not agree on acceptable NCP parameters within the time allowed for upper-layer negotiation.</p> <p>Code 19 takes precedence over code 17 in situations related to address convergence failure.</p>
18	ncp no ncps available	No NCPs were successfully enabled within the time allowed for upper-layer negotiation.
19	ncp addresses failed to converge	<p>An NCP configuration error has prevented NCP negotiation from converging on acceptable addresses. This occurs if the two peers never agree on acceptable NCP addresses within the time allowed for upper-layer negotiation.</p> <ul style="list-style-type: none"> Code 19 with direction 1 is generated if the peer denies address parameters requested by the local NCP. Code 19 with direction 2 is generated if the local NCP denies address parameters requested by the peer. <p>The IPv6 interface identifier is considered an address for the purposes of code 19.</p> <p>Code 19 takes precedence over code 17 in situations related to address convergence failure.</p>
20	ncp negotiation inhibited	<ul style="list-style-type: none"> Code 20 with direction 2 indicates that an upper layer negotiation was inhibited for any enabled NCP because the required network-layer parameters were not available as a result of the authentication stage. Code 20 with direction 1 is never generated; the NCPs are never enabled if there is no non-null local address.

- Related Documentation**
- [L2TP Terminate Reasons on page 266](#)
 - [PPP Terminate Reasons on page 283](#)
 - *disconnect-cause*
 - *l2tp disconnect-cause*

- *show l2tp received-disconnect-cause-summary*

Monitoring L2TP and L2TP Dial-Out

When you have configured L2TP and L2TP dial-out on your E Series router, you can monitor the active tunnels and sessions.



NOTE: All of the commands in this chapter apply to both the LAC and the LNS.

L2TP and L2TP dial-out topics are described in the following sections:

- [Monitoring the Mapping for User Domains and Virtual Routers with AAA on page 427](#)
- [Monitoring Configured Tunnel Groups with AAA on page 430](#)
- [Monitoring Configuration of Tunnel Parameters with AAA on page 432](#)
- [Monitoring Global Configuration Status on E Series Routers on page 433](#)
- [Monitoring Detailed Configuration Information for Specified Destinations on page 435](#)
- [Monitoring Locked Out Destinations on page 437](#)
- [Monitoring Configured L2TP Destination Profiles or Host Profiles on page 438](#)
- [Monitoring Configured and Operational Status of all Destinations on page 440](#)
- [Monitoring Statistics on the Cause of a Session Disconnection on page 441](#)
- [Monitoring Detailed Configuration Information about Specified Sessions on page 442](#)
- [Monitoring Configured and Operational Summary Status on page 443](#)
- [Monitoring Configured Switch Profiles on Router on page 444](#)
- [Monitoring Detailed Configuration Information about Specified Tunnels on page 445](#)
- [Monitoring Configured and Operational Status of All Tunnels on page 448](#)
- [Monitoring Chassis-wide Configuration for L2TP Dial-out on page 449](#)
- [Monitoring Status of Dial-out Sessions on page 454](#)
- [Monitoring Dial-out Targets within the Current VR Context on page 455](#)
- [Monitoring Operational Status within the Current VR Context on page 456](#)

Monitoring the Mapping for User Domains and Virtual Routers with AAA

Purpose Display the mapping between user domains and virtual routers.

Action To display the mapping between user domains and virtual routers:

```
host1#show aaa domain-map
```

```
Domain: lac-tunnel; router-name: lac; ipv6-router-name: default
```

```

Tunnel      Tunnel      Tunnel      Tunnel      Tunnel      Tunnel
Tag         Peer      Source      Type      Medium      Password      Tunnel Id
-----
5          192.168.1.1  <null>      12tp      ipv4        welcome      lac-tunnel

Tunnel      Tunnel      Tunnel      Tunnel      Tunnel
Tag         Client Name Server Name Preference Max Sessions Tunnel RWS
-----
5          lac        boston      5          0          4

Tunnel      Tunnel      Tunnel      Tunnel      Tunnel
Tag         Virtual Router Failover Resync Switch Profile Tx Speed Method
-----
5          <null>      <null>      denver      qos

```

Meaning [Table 89 on page 428](#) lists the **show aaa domain-map** command output fields.

Table 89: show aaa domain-map Output Fields

Field Name	Field Description
Domain	Name of the domain
router-name	Virtual router to which user domain name is mapped
router-mask	IPv4 mask of the local interface
tunnel-group	Name of the tunnel group assigned to the domain map
ipv6-router-name	IPv6 virtual router to which user domain name is mapped
local-interface	Interface information to use on the local (E Series) side of the subscriber's interface
ipv6-local-interface	IPv6 interface information to use on the local (E Series) side of the subscriber's interface
poolname	Local address pool from which the router allocates addresses for this domain
IP hint	IP hint is enabled
strip-domain	Strip domain is enabled
override-username	Single username used for all users from a domain in place of the values received from the remote client

Table 89: show aaa domain-map Output Fields (*continued*)

Field Name	Field Description
override-password	Single password used for all users from a domain in place of the values received from the remote client
Tunnel Tag	Tag that identifies the tunnel
Tunnel Peer	Destination address of the tunnel
Tunnel Source	Source address of the tunnel
Tunnel Type	L2TP
Tunnel Medium	Type of medium for the tunnel; only IPv4 is supported
Tunnel Password	Password for the tunnel
Tunnel Id	ID of the tunnel
Tunnel Client Name	Hostname that the LAC sends to the LNS when communicating to the LNS about the tunnel
Tunnel Server Name	Hostname expected from the peer (the LNS) when during tunnel startup
Tunnel Preference	Preference level for the tunnel
Tunnel Max Sessions	Maximum number of sessions allowed on a tunnel
Tunnel RWS	L2TP receive window size (RWS) for a tunnel on the LAC; displays either the configured value or the default behavior, which is indicated by system chooses (that is, the RWS value is set as 64)
Tunnel Virtual Router	Name of the virtual router to map to the user domain name
Tunnel Failover Resync	L2TP peer resynchronization method
Field descriptions	The actual fields displayed depend on your configuration
Tunnel Switch Profile	Name of the L2TP tunnel switch profile
Tunnel Tx Speed Method	Method that the router uses to calculate the transmit connect speed of the subscriber's access interface: static layer2, dynamic layer2, qos, actual, not set

Related Documentation • [show aaa domain-map](#)

Monitoring Configured Tunnel Groups with AAA

Purpose Display the currently configured tunnel groups.

Action To display information about currently configured tunnel groups:

```
host1#show aaa tunnel-group
```

```
Tunnel Group: boston
```

Tunnel Tag	Tunnel Peer	Tunnel Source	Tunnel Type	Tunnel Medium	Tunnel Password	Tunnel Id
3	192.168.1.1	<null>	l2tp	ipv4	msn	<null>

Tunnel Tag	Tunnel Client Name	Tunnel Server Name	Tunnel Preference	Tunnel Max Sessions	Tunnel RWS
3	msn.del.com	<null>	2000	0	4

Tunnel Tag	Tunnel Virtual Router	Tunnel Failover Resync	Tunnel Switch Profile	Tunnel Tx Speed Method
3	<null>	<null>	sanjose	qos

Meaning [Table 90 on page 430](#) lists the **show aaa tunnel-group** command output fields.

Table 90: show aaa tunnel-group Output Fields

Field Name	Field Description
Domain	Name of the domain
router-name	Virtual router to which user domain name is mapped
router-mask	IPv4 mask of the local interface
tunnel-group	Name of the tunnel group assigned to the domain map
ipv6-router-name	IPv6 virtual router to which user domain name is mapped
local-interface	Interface information to use on the local (E Series) side of the subscriber's interface
ipv6-local-interface	IPv6 interface information to use on the local (E Series) side of the subscriber's interface
poolname	Local address pool from which the router allocates addresses for this domain

Table 90: show aaa tunnel-group Output Fields (*continued*)

Field Name	Field Description
IP hint	IP hint is enabled
strip-domain	Strip domain is enabled
override-username	Single username used for all users from a domain in place of the values received from the remote client
override-password	Single password used for all users from a domain in place of the values received from the remote client
Tunnel Tag	Tag that identifies the tunnel
Tunnel Peer	Destination address of the tunnel
Tunnel Source	Source address of the tunnel
Tunnel Type	L2TP
Tunnel Medium	Type of medium for the tunnel; only IPv4 is supported
Tunnel Password	Password for the tunnel
Tunnel Id	ID of the tunnel
Tunnel Client Name	Hostname that the LAC sends to the LNS when communicating to the LNS about the tunnel
Tunnel Server Name	Hostname expected from the peer (the LNS) when during tunnel startup
Tunnel Preference	Preference level for the tunnel
Tunnel Max Sessions	Maximum number of sessions allowed on a tunnel
Tunnel RWS	L2TP receive window size (RWS) for a tunnel on the LAC; displays either the configured value or the default behavior, which is indicated by system chooses (that is, the RWS value is set as 64)
Tunnel Virtual Router	Name of the virtual router to map to the user domain name
Tunnel Failover Resync	L2TP peer resynchronization method
Field descriptions	The actual fields displayed depend on your configuration
Tunnel Switch Profile	Name of the L2TP tunnel switch profile

Table 90: show aaa tunnel-group Output Fields (*continued*)

Field Name	Field Description
Tunnel Tx Speed Method	Method that the router uses to calculate the transmit connect speed of the subscriber's access interface: static layer2, dynamic layer2, qos, actual, not set

- Related Documentation**
- The information displayed is almost identical to the tunnel information displayed using the **show aaa domain-map** command. See [Monitoring the Mapping for User Domains and Virtual Routers with AAA on page 427](#).
 - `show aaa tunnel-group`

Monitoring Configuration of Tunnel Parameters with AAA

Purpose Display configuration of tunnel parameters used for tunnel definitions.

Action To display the configuration of tunnel parameters used for tunnel definitions:

```
host1#show aaa tunnel-parameters
Tunnel password is 3&92k%b#q4
Tunnel client-name is <NULL>
Tunnel nas-port-method is none
Tunnel switch profile is boston
Tunnel tx-connect-speed-method is qos
Tunnel nas-port ignore disabled
Tunnel nas-port-type ignore disabled
Tunnel assignmentId format is assignmentId
Tunnel calling number format is fixed (stacked)
Tunnel calling number format fallback is fixed
```

Meaning [Table 91 on page 432](#) lists the **show aaa tunnel-parameters** command output fields.

Table 91: show aaa tunnel-parameters Output Fields

Field Name	Field Description
Tunnel password	Default tunnel password
Tunnel client-name	Hostname that the LAC sends to the LNS when communicating about the tunnel
Tunnel nas-port-method	Default NAS port type
Tunnel switch profile is	Name of the default L2TP tunnel switch profile
Tunnel tx-connect-speed-method is	Method that the router uses to calculate the transmit connect speed of the subscriber's access interface: static layer2, dynamic layer2, qos, actual, not set
Tunnel nas-port ignore	Whether the router uses the tunnel peer's NAS-Port [5] attribute; enabled or disabled

Table 91: show aaa tunnel-parameters Output Fields (*continued*)

Field Name	Field Description
Tunnel nas-port-type ignore	Whether the router uses the tunnel peer's NAS-Port-Type [61] attribute; enabled or disabled
Tunnel assignmentId format	Value of the tunnel assignment ID that is passed to PPP/L2TP
Tunnel calling number format	Format configured for L2TP Calling Number AVP 22 generated by the LAC
Tunnel calling number format fallback	Fallback format configured for L2TP Calling Number AVP 22 generated by the LAC

Related Documentation

- *show aaa tunnel-parameters*

Monitoring Global Configuration Status on E Series Routers

Purpose Display the global configuration and status for L2TP on E Series routers, including switched sessions.

Action To display the global configuration and status for L2TP on E Series routers, including switched sessions:

```

host1#show l2tp
Configuration
L2TP administrative state is enabled
Dynamic interface destruct timeout is 600 seconds
Data packet checksums are disabled
Receive data sequencing is not ignored
Tunnel switching is disabled
Retransmission retries for established tunnels is 5
Retransmission retries for not-established tunnels is 5
Tunnel idle timeout is 60 seconds
Failover within a preference level is disabled
Weighted load balancing is disabled
Tunnel authentication challenge is enabled
Calling number avp is enabled
L2TP fast path hello processing during ISSU is enabled
L2TP Call Disconnect on Local pool exhaustion is disabled
Reject remote transmit address change is enabled for ip address
Ignore remote transmit address change is disabled
Disconnect-cause avp generation is enabled
Default receive window size is system chooses
Rx speed avp when equal is enabled
Destination lockout timeout is 300 seconds
Destination lockout test is disabled
Failover resync is silent-failover
Sub-interfaces      total      active      failed      auth-errors
Destinations        0          0          0          n/a
Tunnels             0          0          0          0
Sessions            0          0          0          n/a
Switched-sessions  0          0          0          n/a

```

Meaning [Table 92 on page 434](#) lists the **show l2tp** command output fields.

Table 92: show l2tp Output Fields

Field Name	Field Description
Configuration	Configuration and status for L2TP on E Series routers, including switched sessions
L2TP administrative state	Status of L2TP on the router; enabled or disabled
Dynamic interface destruct timeout	Number of seconds that the router maintains dynamic destinations, tunnels, and sessions after they have terminated
Data packet checksums	Status of checking data integrity via UDP; enabled or disabled
Receive data sequencing	Whether the router processes or ignores sequence numbers in incoming data packets
Tunnel switching	Enabled or disabled
Retransmission retries for established tunnels	Number of retries configured for established tunnels
Retransmission retries for not-established tunnels	Number of retries configured for tunnels not established
Tunnel idle timeout	Length of the tunnel idle timeout, in seconds
Failover within a preference level	Enabled or disabled
Weighted load balancing	Enabled or disabled
Tunnel authentication challenge	Enabled or disabled
Calling number avp	Whether the E Series LAC sends Calling-Station-Id and Called-Station-Id AVPs in ICRQ packets, enabled or disabled
L2TP fast path hello processing during ISSU	Status of handling L2TP hello packets in fast path during unified ISSU; enabled or disabled.
L2TP Call Disconnect on Local pool exhaustion	Status of sending an L2TP CDN message to LAC on IPv4 local pool exhaustion; enabled or disabled.
Reject remote transmit address change	Enabled or disabled for IP address, UDP port, or both
Ignore remote transmit address change	Enabled or disabled for IP address, UDP port, or both

Table 92: show l2tp Output Fields (*continued*)

Field Name	Field Description
Disconnect-cause avp generation	Enabled or disabled
Default receive window size	Default L2TP RWS for a tunnel on both the LAC and the LNS; displays either the configured value or the default behavior, indicated by system chooses (that is, the RWS value is set as 64)
Rx speed avp when equal	Enabled or disabled
Destination lockout timeout	Number of seconds that L2TP destinations remain in the lockout state after they become unavailable
Destination lockout test	Status of the L2TP destination lockout test, enabled or disabled
Failover resync	Global L2TP peer resynchronization configuration
Sub-interfaces	Sub-interface information about L2TP
total	Number of destinations, tunnels, and sessions that the router created
active	Number of operational destinations, tunnels, and sessions
failed	Number of requests that did not reach an operational state
auth-errors	Number of requests that failed because the tunnel password was invalid

Related Documentation

- *show l2tp*

Monitoring Detailed Configuration Information for Specified Destinations

Purpose Display detailed configuration information about specified destinations.

Action To display detailed configuration information about specified destinations:

To display information about a specific destination:

```
host1#show l2tp destination ip 172.31.1.98
```

```
L2TP destination 1 is Up with 5 active tunnels and 64 active sessions
```

To display information about all destinations:

```
host1#show l2tp destination detail 1
```

```
L2TP destination 1 is Up with 5 active tunnels and 64 active sessions
Configuration
```

```

Administrative state is enabled
SNMP traps are enabled
Destination address
Transport ipUdp
Virtual router default
Local address 192.168.1.230, peer address 172.31.1.98
Destination status
Effective administrative state is enabled
Sub-interfaces total active failed auth-errors
Tunnels      5      5      0      0
Sessions     64     64      0     n/a
Statistics   packets  octets   discards  errors
Control rx   69             3251       2         0
Control tx   195            23939      0         0
Data rx      68383456      68383456  0         0
Data tx      68383456      68383456  0         0

```

Meaning [Table 93 on page 436](#) lists the **show l2tp destination** command output fields.

Table 93: show l2tp destination Output Fields

Field Name	Field Description
Configuration	Configured status of the destination
Administrative state	Administrative status of the destination: <ul style="list-style-type: none"> enabled—No restrictions on creation and operation of sessions and tunnels for this destination disabled—Router disabled existing sessions and tunnels and will not create new sessions or tunnels for this destination drain—Router will not create new sessions or tunnels for this destination
SNMP traps	Whether or not the router sends traps to SNMP for operational state changes
Destination address	Address information for the specified destination
Transport	Method used to transfer traffic
Virtual	Name of the virtual router on which the tunnel is configured
Local and peer addresses	Addresses of the local and remote interfaces
Destination status	Effective administrative state—The more restrictive of the router and destination administrative states. This setting, rather than the administrative state of the destination, determines whether the router can create new sessions or tunnels and whether the sessions or tunnels are disabled for this destination.
Sub-interfaces	Sub-interface information about the L2TP destination

Table 93: show l2tp destination Output Fields (*continued*)

Field Name	Field Description
total	Number of sessions or tunnels that the router created for this destination
active	Number of operational sessions or tunnels for this destination
failed	Number of requests that did not reach an operational state for this destination
auth-errors	Number of requests that failed because the tunnel password was invalid for this destination
Statistics	Information about the traffic sent and received

Related Documentation

- *show l2tp destination*

Monitoring Locked Out Destinations

Purpose Display information about the L2TP destinations that are currently locked out.

Action To display information about the L2TP destinations that are currently locked out:

```
host1#show l2tp destination lockout
L2TP destination 36 is waiting for lockout timeout (45 seconds remaining)
L2TP destination 54 is waiting for lockout test start
L2TP destination 76 is waiting for lockout test complete
3 L2TP lockout destinations found
```

Meaning Table 94 on page 437 lists the **show l2tp destination lockout** command output fields.

Table 94: show l2tp destination lockout Output Fields

Field Name	Field Description
L2TP destination waiting	Name of destination and its lockout status. The status indicates whether the destination is waiting for the lockout timeout to expire (and how much time is left), or waiting for the lockout test to start or finish
L2TP lockout destinations found	Number of destinations that are currently in lockout state

Related Documentation

- *show l2tp destination lockout*

Monitoring Configured L2TP Destination Profiles or Host Profiles

Purpose Display either a list of configured Layer 2 Tunneling Protocol (L2TP) destination profiles or the host profiles defined in a particular profile.

If a nondefault L2TP receive window size (RWS) is configured for a particular host profile, the command displays the RWS setting as an attribute of that host profile.

Action To display either a list of configured L2TP destination profiles or the host profiles defined in a particular profile:

```
host1#show l2tp destination profile
L2TP destination profile westford
1 L2TP destination profile found
```

If a nondefault L2TP RWS is configured for a particular host profile, to display the RWS setting as an attribute of that host profile:

```
host1#show l2tp destination profile westford
L2TP destination profile westford
Configuration
  Destination address
  Transport ipUdp
  Virtual router lns
  Peer address 192.168.1.99
  Destination profile maximum sessions is 5000
Current session count in group-A is 14, max-sessions configured is 3400
Current session count in group-B is 2, max-sessions configured is 4600
Statistics
  Destination profile current session count is 30
Host profile attributes
  Remote host is remhost22.xyz.com
  Configuration
    Tunnel password is 23erf5
    Interface profile is ebcints
    Bundled group id is 1
    Bundled group id override is enabled
    Maximum sessions is 400
    Failover resync is failover-protocol
    Sessions-limit-group is group-A
  Statistics
    Current session count is 14
  Remote host is asciitext
  Configuration
    Bundled group id is 0
    Tunnel password is 222
    Interface profile is ascints
    Default upper binding type mlppp
    Maximum sessions is 250
    Failover resync is failover-protocol
    Sessions-limit-group is group-B
  Statistics
    Current session count is 2
  Remote host is mexico
  Configuration
    Local ip address is 10.10.2.2
    Proxy lcp is disabled
    Proxy authenticate is enabled
    mlppp upper binding type
```

```

Disconnect-cause avp is enabled
Receive window size is 4
Maximum sessions is 500
Failover resync is failover-protocol
Statistics
Current session count is 14
Remote host is LAC
Configuration
Tunnel password is TunnelPass
Local host name is LNS
Local ip address is 46.1.1.2
Session virtual router is vr1
Disconnect-cause avp is enabled
Tunnels are single-shot
Override out-of-resource-result-code is enabled
Statistics
Current session count is 0
5 L2TP host profiles found

```

Meaning Table 95 on page 439 lists the **show l2tp destination profile** command output fields.

Table 95: show l2tp destination profile Output Fields

Field Name	Field Description
Transport	Method used to transfer traffic
Virtual router	Name of the virtual router
Peer address	IP address of the L2TP access concentrator (LAC)
Destination profile maximum sessions	Maximum number of sessions allowed for the destination profile
Current session count in group-A	Number of current sessions in group-A
Current session count in group-B	Number of current sessions in group-B
Destination profile current session count	Number of current sessions for the destination profile
Host profile attributes	Host profile attributes of the L2TP destination
Remote host	Name of the remote host
Local host name	Name of the local host
Local ip address	IP address of the local host
Bundled group id	Identifier for bundled sessions
Bundled group id override	Status of the bundled group ID override: enabled or disabled
Tunnel password	Password for the tunnel

Table 95: show l2tp destination profile Output Fields (*continued*)

Field Name	Field Description
Interface profile	Name of the host profile
Default upper binding type	The default upper binding type: mlpp
Proxy lcp	Status of the proxy LCP for the remote host
mlppp upper binding type	Default upper binding type
Proxy authenticate	The status of the proxy authentication: enabled or disabled
Disconnect-cause avp	Status of the disconnect-cause attribute-value pair (AVP): enabled or disabled
Tunnels are single-shot	Indicates that single-shot tunnels are configured for this host profile
Receive window size	Number of packets that the peer can transmit without receiving an acknowledgment from the router
Maximum sessions	Maximum number of sessions allowed for the host profile
Failover resync	L2TP peer resynchronization method for the host profile
Override out-of-resource-result-code	State of the out-of-resource-result-code override: enabled or disabled
Current session count	Number of current sessions for the host profile
Sessions-limit-group	Name of the sessions limit group
Session virtual router	Name of the virtual router assigned to the L2TP host profile for mapping the PPP session to the TSMs.

- Related Documentation**
- [Configuring an LNS on page 366](#)
 - [Configuring an L2TP Destination Profile to Enable IPsec Support for L2TP Tunnels](#)
 - [Configuring Single-Shot L2TP/IPsec Tunnels](#)
 - [show l2tp destination profile](#)

Monitoring Configured and Operational Status of all Destinations

Purpose Display summary of the configured and operational status of all L2TP destinations.

Action To display a summary of the configured and operational status of all L2TP destinations:

```
host1#show l2tp destination summary
Administrative status   enabled      drain        disabled
                        0            0            0
Operational status     up           down         lower-down   not-present
                        0            0            0            0
```

Meaning [Table 96 on page 441](#) lists the **show l2tp destination summary** command output fields.

Table 96: show l2tp destination summary Output Fields

Field Name	Field Description
Administrative status	<p>Administrative status of the L2TP destination:</p> <ul style="list-style-type: none"> enabled—No restrictions on creation and operation of sessions and tunnels for this destination drain—Router will not create new sessions or tunnels for this destination disabled—Router disabled existing sessions and tunnels and will not create new sessions or tunnels for this destination
Operational status	<p>Operational status of the L2TP destination:</p> <ul style="list-style-type: none"> up—Destination is available for tunnels down—Destination is not available for tunnels lower-down—Underlying transport is unavailable; for example, you removed the virtual router not-present—Hardware supporting the destination is unavailable; for example, you removed a required line module

Related Documentation

- show l2tp destination*

Monitoring Statistics on the Cause of a Session Disconnection

Purpose Display statistics for all information the LAC receives from an LNS about the cause of an L2TP session disconnection.

Action To display statistics for all information the LAC receives from an LNS about the cause of an L2TP session disconnection.

```
host1# show l2tp received-disconnect-cause-summary
Disconnect Cause (Code)      Global    Peer      Local
-----
no info (0)                  0         0         0
admin disconnect (1)         0         0         0
renegotiation disabled (2)   0         0         0
normal disconnect (3)        0         0         0
compulsory encryption refused (4) 0         0         0
lcp failed to converge (5)    0         0         0
lcp peer silent (6)          0         0         0
lcp magic number error (7)    0         0         0
```

lcp keepalive failure (8)	0	0	0
lcp mlppp endpoint discriminator mismatch (9)	0	0	0
lcp mlppp peer mrru not valid (10)	0	0	0
lcp mlppp peer ssn invalid (11)	0	0	0
lcp callback refused (12)	0	0	0
authenticate timed out (13)	0	0	0
authenticate mlppp name mismatch (14)	0	0	0
authenticate protocol refused (15)	0	0	0
authenticate failure (16)	0	0	0
ncp no negotiation completed (17)	0	0	0
ncp no ncps available (18)	0	0	0
ncp addresses failed to converge (19)	0	0	0
ncp negotiation inhibited (20)	0	0	0

Meaning [Table 97 on page 442](#) lists the `show l2tp received-disconnect-cause-summary` command details.

Table 97: show l2tp received-disconnect-cause-summary Output Fields

Field Name	Field Description
show l2tp received-disconnect-cause-summary	Display statistics for all information the LAC receives from an LNS about the cause of an L2TP session disconnection.

Related Documentation

- *show l2tp received-disconnect-cause-summary*

Monitoring Detailed Configuration Information about Specified Sessions

Purpose Display detailed configuration information about specified sessions.

Action To display detailed configuration information about specified sessions:

To display L2TP session:

```
host1#show l2tp session
L2TP session 1/1/1 is Up
1 L2TP session found
```

To display L2TP session details:

```
host1#show l2tp session detail
L2TP session 1/1/1 is Up
Configuration
  Administrative state is enabled
  SNMP traps are enabled
Session status
  Effective administrative state is enabled
  State is established
  Local session id is 25959, peer session id is 2
Statistics packets octets discards errors
Data rx 7      237    1      0
Data tx 6      160    0      0
```

```
Session operational configuration
  User name is 't1.s1@local'
```

```

Tunneling PPP interface atm 0/0.1
Call type is lacIncoming
Call serial number is 0
Bearer type is none
Framing type is none
Proxy LCP was provided
Authentication method was chap
Tunnel switch profile is chicago

```

Meaning Table 98 on page 443 lists the **show l2tp session** command output fields.

Table 98: show l2tp session Output Fields

Field Name	Field Description
Configuration	Configured status of the session
Administrative state	Administrative status of the destination: <ul style="list-style-type: none"> • enabled—No restrictions on the operation of this session • disabled—Router terminated this session
SNMP traps	Whether or not the router sends traps to Simple Network Management Protocol (SNMP) for operational state changes
Session status	Session status of the destination
Effective administrative state	Most restrictive of the following administrative states: router, destination, tunnel, and session. This setting, rather than the administrative state of the session, determines whether the router can maintain this session or not.
State	Status of the session: idle, connecting, established, or disconnecting
Local and peer session id	Names the router uses to identify the session locally and remotely
Statistics	Information about the traffic for this session
Session operational configuration	Information received from the peer when the session was created

Related Documentation

- *show l2tp session*

Monitoring Configured and Operational Summary Status

Purpose Display a summary of the configured and operational status of all L2TP sessions.

Action To display a summary of the configured and operational status of all L2TP sessions:

```
host1# show l2tp session summary
Administrative status  enabled    disabled
                      64        0
Operational status    up        down    lower-down    not-present
                      64        0        0            0
```

Meaning [Table 99 on page 444](#) lists the **show l2tp session summary** command output fields.

Table 99: show l2tp session summary Output Fields

Field Name	Field Description
Administrative status:	Administrative status of the session: <ul style="list-style-type: none"> • enabled—No restrictions on the creation of sessions • disabled—Router disabled these sessions
Operational status:	Operational status of the session: <ul style="list-style-type: none"> • up—Session is available • down—Session is unavailable • lower-down—Session is unavailable because the tunnel supporting it is inaccessible • not-present—Session is unavailable because the hardware (such as a line module) supporting it is inaccessible

Related Documentation

- *show l2tp session summary*

Monitoring Configured Switch Profiles on Router

Purpose Display information about the L2TP switch profiles configured on the router.

Action To display only the names of the L2TP tunnel switch profiles configured on the router:

```
host1#show l2tp switch-profile
L2TP tunnel switch profile concord
L2TP tunnel switch profile myProfile
2 L2TP tunnel switch profiles found
```

To display information about the settings in a particular L2TP tunnel switch profile:

```
host1#show l2tp switch-profile concord
L2TP tunnel switch profile concord
  AVP bearer type action is relay
  AVP calling number action is relay
  AVP Cisco nas port info action is relay
```

Meaning [Table 100 on page 445](#) lists the **show l2tp switch-profile** command output fields.

Table 100: show l2tp switch-profile Output Fields

Field Name	Field Description
L2TP tunnel switch profile	Name of the L2TP tunnel switch profile
AVP <i>actionType</i> action is	Indicates the tunnel switching behavior or action type (for example, relay) configured for the specified L2TP AVP type

Related Documentation

- *show l2tp switch-profile*

Monitoring Detailed Configuration Information about Specified Tunnels

Purpose Display detailed configuration information about specified tunnels.

Action To display detailed configuration information about specified tunnel by IP address:



NOTE: For L2TP tunnels configured with output IPv4 or IPv6 policy lists, the output of the `show ip interface tunnel l2tp:tunnel-name` and the `show ipv6 interface tunnel l2tp:tunnel-name` commands display only the forwarded packets and bytes fields, and the dropped packets and bytes fields in the rate-limit-profile section for policies with hierarchical parent groups under the IP policy output or IPv6 policy output headings, respectively, when scheduler profile–based computation of service session accounting is enabled. In such a case, the committed, conformed, exceeded, saturated, and unconditional packets and bytes fields are not displayed in the rate-limit-profile section in the output of these commands for policies with hierarchical parent groups.

```
host1#show l2tp tunnel virtual router default ip 172.31.1.98
L2TP tunnel 1/xyz is Up with 13 active sessions
L2TP tunnel 1/aol.com is Up with 13 active sessions
L2TP tunnel 1/isp.com is Up with 13 active sessions
L2TP tunnel 1/msn.com is Up with 13 active sessions
L2TP tunnel 1/mv.com is Up with 12 active sessions
5 L2TP tunnels found
```

To display detailed configuration information about specified tunnel:

```
host1#show l2tp tunnel detail 1/xyz
L2TP tunnel 1/xyz is Up with 13 active sessions
Configuration
  Administrative state is enabled
  SNMP traps are enabled
Tunnel address
  Transport ipUdp
  Virtual router default
  Local address 192.168.1.230, peer address 172.31.1.98
```

```

Local UDP port 1701, peer UDP port: 1701
Tunnel status
Effective administrative state is enabled
State is established
Local tunnel id is 14529, peer tunnel id is 34
Host profile is none
Tunnel is Up for: 12 days, 8 hours, 24 minutes, 23 seconds
Sub-interfaces      total  active  failed
Sessions            13      13      0
Statistics  packets  octets    discards  errors
Control rx   14      683          0         0
Control tx   41     4666          0         0
Data rx      67900944 67900944    0         0
Data tx      67900944 67900944    0         0
Control channel statistics
Receive window size = 4
Receive ZLB = 17
Receive out-of-sequence = 0
Receive out-of-window = 0
Transmit window size = 4
Transmit ZLB = 12
Transmit queue depth = 0
Retransmissions = 8
Tunnel operational configuration
Peer host name is 'Juniper-POS'
Peer vendor name is 'XYZ, Inc.'
Peer protocol version is 1.1
Peer firmware revision is 0x1120
Peer bearer capabilities are digital and analog
Peer framing capabilities are sync and async

```

Meaning [Table 101 on page 446](#) lists the **show l2tp tunnel** command output fields.

Table 101: show l2tp tunnel Output Fields

Field Name	Field Description
Configuration	Configured status of the tunnel enabled.
Administrative state	Administrative status of the enabled tunnel: <ul style="list-style-type: none"> enabled—No restrictions on creation and operation of sessions for this tunnel disabled—Router disabled existing sessions and will not create new sessions on this tunnel drain—Router will not create new sessions on this tunnel
SNMP traps	Whether or not the router sends traps to SNMP for operational state changes.
Tunnel address	Tunnel address information.
Transport	Method used to transfer traffic.
Virtual router	Name of the virtual router on which the tunnel is configured.

Table 101: show l2tp tunnel Output Fields (*continued*)

Field Name	Field Description
Local and peer addresses	IP addresses of the local and remote ends of the tunnel. If the router is set up to ignore address and port changes in SCCRP packets, both the transmit and receive addresses are listed for the peer.
Local and peer UDP ports	UDP ports for the local and remote ends of the tunnel. If the router is set up to accept address and port changes in SCCRP packets, both the transmit and receive UDP ports are listed for the peer.
Tunnel status	Tunnel status information.
Effective administrative state	Most restrictive of the following administrative states: E Series router, destination, and tunnel. This setting, rather than the administrative state of the tunnel, determines whether the router can create new sessions on a tunnel or whether the sessions on a tunnel are disabled or not.
State	Status of the enabled tunnel: <ul style="list-style-type: none"> • idle • connecting • established • disconnecting
Local and peer tunnel id	Names the router used to identify the tunnel locally and remotely.
Host profile	Name of the L2TP host profile, if it is configured. Otherwise, the label “none” is displayed to specify that a host profile is not enabled.
Tunnel is Up for	Duration for which the tunnel is operationally up, which is denoted in terms of days, hours, minutes, and seconds.
Sub-interfaces:	Sub-interface information for the enabled tunnel: <ul style="list-style-type: none"> • total—Number of sessions that the router has created on this tunnel • active—Number of operational sessions on the tunnel • failed—Number of requests that did not reach an operational state
Statistics	Information about the traffic sent and received.
Control channel statistics	Tunnel control channel information.

Table 101: show l2tp tunnel Output Fields (*continued*)

Field Name	Field Description
Receive window size	Number of packets that the peer can transmit without receiving an acknowledgment from the router.
Receive ZLB	Number of acknowledgments that the router has received from the peer.
Receive out-of-sequence	Number of received control packets that were out of order.
Receive out-of-window	Number of packets that arrived at the router outside the receiving window.
Transmit window size	Number of packets that the router can transmit before receiving an acknowledgment from the peer.
Transmit ZLB	Number of acknowledgments that the router has sent to the peer.
Transmit queue depth	Number of packets that the router is waiting to send to the peer, plus the number of packets for which the peer has not yet acknowledged receipt.
Tunnel operation configuration	Information received from the peer when the tunnel was created.

Related Documentation

- [show l2tp tunnel](#)

Monitoring Configured and Operational Status of All Tunnels

Purpose Display a summary of the configured and operational status of all L2TP tunnels.

Action To display a summary of the configured and operational status of all L2TP tunnels:
 host1#show l2tp tunnel summary

```

Administrative status  enabled    drain    disabled
                    5          0         0
Operational status    up        down    lower-down  not-present
                    5          0         0           0

```

Meaning [Table 102 on page 449](#) lists the **show l2tp tunnel summary** command output fields.

Table 102: show l2tp tunnel summary Output Fields

Field Name	Field Description
Administrative status	Administrative status of all tunnels: <ul style="list-style-type: none"> • enabled—No restrictions on the creation and operation of sessions for this tunnel • drain—Router will not create new sessions for this tunnel • disabled—Router disabled existing sessions and will not create new sessions for this tunnel
Operational status	Operational status of all tunnels: <ul style="list-style-type: none"> • up—Tunnel is available • down—Tunnel is unavailable • lower-down—Tunnel is unavailable because the destination supporting it is inaccessible • not-present—Tunnel is unavailable because the hardware (such as a line module) supporting the tunnel is inaccessible

Related Documentation

- *show l2tp tunnel summary*

Monitoring Chassis-wide Configuration for L2TP Dial-out

Purpose To display the chassis-wide configuration, operational state, and statistics for L2TP dial-out.

This command displays aspects of the dial-out state machine and details about the dial-out routes themselves. This section presents sample output. The actual output on your router may differ significantly.

Action To display chassis-wide configuration, operational state, and statistics for L2TP dial-out:

```
host1#show l2tp dial-out
Operational status: inService
Connecting timer value: 30 seconds
Dormant timer value: 300 seconds
```

To display detailed chassis-wide configuration information:

```
host1#show l2tp dial-out detail
Dial-out Chassis Configuration and Operational Status
Chassis operational status : inService
Dormant timeout           : 30 seconds
Connecting timeout        : 30 seconds

Dial-out Chassis Statistics
Current sessions:                0
Maximum sessions:                0
Current sessions in the process of connecting: 0
Maximum sessions connecting at one time: 0
Current sessions pending:        0
```

```

Maximum sessions pending:                0
Current targets inhibited:                0
Maximum targets inhibited:                0
Authentication grant for nonexistent session: 0
Authentication deny for nonexistent session: 0

Dial-out Virtual router statistics
Virtual routers active:                  0
Virtual routers created:                  0
Virtual routers removed:                  0
Virtual routers in init-pending state:    0
Virtual routers in init-failed state:     0
Virtual routers in down state:            0
Virtual routers in in-service state:      0
IP Discarded trigger frames:              0
Trigger frames received for unknown route: 0
Sessions in dormant state:                0
Sessions in pending state:                0
Sessions in authenticating state:          0
Sessions in connecting state:              0
Sessions in in-service state:              0
Sessions in inhibited state:              0
Sessions in post-inhibited state:         0
Sessions in failed state:                 0

Dial-out target statistics
Targets active:                          0
Targets created:                          0
Targets removed:                          0
Targets in down state:                    0
Targets in inhibited state:                0
Targets in in-service state:              0
Triggers discarded:                       0

Dial-out session statistics
Sessions active:                          0
Sessions created:                          0
Sessions removed:                          0
Sessions reset:                            0
Triggers received:                        0
Triggers enqueued:                        0
Triggers discarded:                       0
Triggers forwarded:                       0
Triggers max enqueued:                    0
Authentication requests:                  0
No resources for authentication:           0
Authentication grants:                    0
Authentication Denies:                    0
Dial-outs requested:                      0
Dial-outs rejected:                       0
Dial-outs established:                    0
Dial-outs timed out:                      0
Dial-outs torn down:                      0

```

To display summary information for chassis-wide configuration:

```

host1#show l2tp dial-out summary
Virtual routers in init pending state :      0
Virtual routers in init failed state  :      0
Virtual routers in down state         :      0
Virtual routers in inService state    :      0
Targets in down state                  :      0
Targets in inhibited state             :      0

```

```

Targets in inService state           :           0
Sessions in dormant state            :           0
Sessions in pending state             :           0
Sessions in authenticating state      :           0
Sessions in connecting state          :           0
Sessions in inService state           :           0
Sessions in inhibited state           :           0
Sessions in postInhibited state       :           0
Sessions in failed state              :           0

```

To display information about the operational or administrative state:

```
host1#show l2tp dial-out state inService
```

Meaning [Table 103 on page 451](#) lists the **show l2tp dial-out** command output fields.

Table 103: show l2tp dial-out Output Fields

Field Name	Field Description
Operational status	Current operational status of the chassis
Connecting timer value	Configuration of the connecting timeout
Dormant timer value	Configuration of the dormant timeout
Dial-out Chassis Statistics	Statistics at the chassis level
Current sessions	Total number of session currently active on the chassis
Maximum sessions	Highest value of current sessions recorded on the chassis since the last router restart
Current sessions in the process of connecting	Sessions currently in the connecting state
Maximum sessions connecting at one time	Highest number of sessions recorded on the chassis at the same time since the last router restart
Current sessions pending	Sessions in the pending state
Maximum sessions pending	Highest number of sessions recorded in the pending state since the last router restart
Current targets inhibited	Targets currently in the inhibited state
Maximum targets inhibited	Highest value of targets recorded in the inhibited state since the last router restart
Authentication grant for nonexistent session	Number of authentication requests granted to nonexistent sessions
Authentication deny for nonexistent session	Number of authentication requests denied to nonexistent sessions

Table 103: show l2tp dial-out Output Fields (*continued*)

Field Name	Field Description
Dial-out Virtual router statistics	Statistics at the virtual router level
Virtual routers active	VRs in use by the state machine
Virtual routers created	VRs that have been used by the state machine
Virtual routers removed	VRs no longer used by the state machine
Virtual routers in init-pending state	VRs in the initializationPending state
Virtual routers in init-failed state	VRs in the initializationFailed state
Virtual routers in down state	VRs in the down state
Virtual routers in in-service state	VRs in the inService state
IP Discarded trigger frames	Trigger frames that IP discarded
Trigger frames received for unknown route	Trigger frames received for an unknown route
Sessions in dormant state	Sessions on the VR that are in the dormant state
Sessions in pending state	Sessions on the VR that are in the pending state
Sessions in authenticating state	Sessions on the VR that are in the authenticating state
Sessions in connecting state	Sessions on the VR that are in the connecting state
Sessions in in-service state	Sessions on the VR that are in the inService state
Sessions in inhibited state	Sessions on the VR that are in the inhibited state
Sessions in post-inhibited state	Sessions on the VR that are in the postInhibited state
Sessions in failed state	Sessions on the VR that are in the failed state
Dial-out target statistics	Statistics at the route target level
Targets active	Current active targets
Targets created	All targets created
Targets removed	Targets deleted
Targets in down state	Targets in the down state

Table 103: show l2tp dial-out Output Fields (*continued*)

Field Name	Field Description
Targets in inhibited state	Targets in the inhibited state
Targets in in-service state	Targets in the inService state
Triggers discarded	Trigger packets discarded
Dial-out session statistics	Statistics at the session level
Sessions active	Currently active sessions
Sessions created	All sessions created
Sessions removed	Sessions deleted
Sessions reset	Sessions reset using the l2tp dial-out session reset command
Triggers received	Triggers received for dial-out sessions
Triggers enqueued	Triggers that have been put into the queue
Triggers discarded	Trigger packets discarded
Triggers forwarded	Trigger packets forwarded
Triggers max enqueued	Maximum number of triggers that have been enqueued simultaneously since the last router reset
Authentication requests	Authentication requests received
No resources for authentication	Authentication requests not processed because of insufficient resources
Authentication grants	Authentication requests granted
Authentication Denies	Authentication requests denied
Dial-outs requested	Outgoing calls requested for sessions
Dial-outs rejected	Outgoing call requests that were rejected
Dial-outs established	Successful outgoing calls before the connecting timer expired
Dial-outs timed out	Number of times the connecting timer expired
Dial-outs torn down	Successful outgoing calls that were terminated

- Related Documentation**
- [L2TP Dial-Out Operational States on page 411](#)
 - *show l2tp dial-out*
 - *show l2tp dial-out virtual-router*

Monitoring Status of Dial-out Sessions

Purpose Display the status of dial-out sessions.

This command displays aspects of the dial-out state machine and details about the dial-out routes themselves. This section presents sample output. The actual output on your router may differ significantly.

Action To display all sessions within the current virtual router context:

```
host1#show l2tp dial-out session
Session           Status
-----
10.10.1.1         connected
10.10.2.1         dormant
```

To display detailed information about a particular session, specify the trigger IP address for the session:

```
host1#show l2tp dial-out session 10.1.1.1
Session 10.1.1.1
Operational status: dormant
```

To display aggregate counts for dial-out sessions in each of the possible operational and administrative states:

```
host1#show l2tp dial-out session summary
```

To display detailed configuration, state, and statistics:

```
host1#show l2tp dial-out session detail
```

To display information about the operational or administrative state:

```
host1#show l2tp dial-out session state connecting
```

To display dial-out information across all virtual routers

```
host1#show l2tp dial-out session allVirtualRouters
```



NOTE: The level of a user's permission determines the use of the **allVirtualRouters** option. For example, if you have permission to view only the current virtual router, then that is all that is displayed when you enter a command.

Meaning [Table 104 on page 455](#) lists the **show l2tp dial-out session** command output fields.

Table 104: show l2tp dial-out session Output Fields

Field Name	Field Description
Session	IP address of the session
Status	Current status of the session
Operational status	Current operational status of session

- Related Documentation**
- [L2TP Dial-Out Operational States on page 411](#)
 - *show l2tp dial-out session*

Monitoring Dial-out Targets within the Current VR Context

Purpose Display configured dial-out targets within the current virtual router context.

This command displays aspects of the dial-out state machine and details about the dial-out routes themselves. This section presents sample output. The actual output on your router may differ significantly.

Action To display general information for all targets within the virtual router:

```
host1:dialout#show l2tp dial-out target
Target      Status      Active Sessions
-----
10.10.1.1/16 up          14
10.1.1.0/24 up          10
```

To display detailed information about a particular target, specify the target IP address and mask:

```
host1:dialout#show l2tp dial-out target 10.1.1.0/24
Target 10.1.1.0/24
Operational status: up
Active sessions: 10
Total triggers: 127
Failed sessions: 2
Connected sessions: 8
```

To display aggregate counts for targets in each of the possible operational and administrative states:

```
host1:dialout#show l2tp dial-out target summary
```

To display detailed configuration, state, and statistics:

```
host1:dialout#show l2tp dial-out target detail
```

To display information about the operational or administrative state:

```
host1:dialout#show l2tp dial-out target state inService
```

To displays dial-out information across all virtual routers:

```
host1:dialout#show l2tp dial-out target allVirtualRouters
```



NOTE: The level of a user's permission determines the use of the **allVirtualRouters** option. For example, if you have permission to view only the current virtual router, then that is all that is displayed when you enter a command.

Meaning Table 105 on page 456 lists the **show l2tp dial-out target** command output fields.

Table 105: show l2tp dial-out target Output Fields

Field Name	Field Description
Target	Address of the target
Status	Status of the connection to the target
Active Sessions	Currently active session to the target
Total triggers	Trigger packets received for the target
Failed sessions	Sessions that are currently in the failed state
Connected sessions	Sessions that are currently in the connected state

- Related Documentation**
- [L2TP Dial-Out Operational States on page 411](#)
 - *show l2tp dial-out target*

Monitoring Operational Status within the Current VR Context

Purpose Display dial-out state machine operational status and statistics within the current VR context.

This command displays aspects of the dial-out state machine and details about the dial-out routes themselves. This section presents sample output. The actual output on your router may differ significantly.

Action To display dial-out state machine operational status and statistics within the current VR context:

```
host1#show l2tp dial-out virtual-router
Dial-out Virtual Router Configuration and Operational Status
Virtual router host1:
Virtual router operational status: inService
Maximum trigger buffers per session: 0
```


To display aggregate counts for dial-out state machines in each of the possible operational and administrative states:

```
host1: dialout#show l2tp dial-out virtual-router summary
```

To display detailed configuration, state, and statistics:

```
host1: dialout#show l2tp dial-out virtual-router detail
```

To display information about the operational or administrative state:

```
host1: dialout#show l2tp dial-out virtual-router state down
```

To displays dial-out information across all virtual routers:

```
host1: dialout#show l2tp dial-out virtual-router allVirtualRouters
```



NOTE: The level of a user's permission determines the use of the **allVirtualRouters** option. For example, if you have permission to view only the current virtual router, then that is all that is displayed when you enter a command.

Meaning [Table 106 on page 457](#) lists the **show l2tp dial-out virtual-router** command output fields.

Table 106: show l2tp dial-out virtual-router Output Fields

Field Name	Field Description
Virtual router	Name of VR
Virtual router operational status	Operational status of the VR
Maximum trigger buffers per session	Maximum number of trigger packets held in buffer while the dial-out session is being established

- Related Documentation**
- [L2TP Dial-Out Operational States on page 411](#)
 - *show l2tp dial-out virtual-router*

PART 4

Managing DHCP

- [DHCP Overview on page 461](#)
- [DHCP Local Server Overview on page 471](#)
- [Configuring DHCP Local Server on page 485](#)
- [Configuring DHCP Relay on page 511](#)
- [Configuring the DHCP External Server Application on page 545](#)
- [Monitoring and Troubleshooting DHCP on page 559](#)

CHAPTER 18

DHCP Overview

Dynamic Host Configuration Protocol (DHCP) provides a mechanism through which computers using TCP/IP can obtain protocol configuration parameters automatically from a DHCP server on the network.

The following sections provide overview information for the E Series router DHCP support:

- [DHCP Overview Information on page 461](#)
- [DHCP Platform Considerations on page 463](#)
- [DHCP References on page 463](#)
- [Configuring the DHCP Access Model on page 464](#)
- [Configuring DHCP Proxy Clients on page 464](#)
- [Logging DHCP Packet Information on page 465](#)
- [Viewing and Deleting DHCP Client Bindings on page 466](#)
- [DHCP Client Bindings and Duplicate MAC Addresses for Subinterfaces Overview on page 469](#)

DHCP Overview Information

The most important configuration parameter carried by DHCP is the IP address. A computer must be initially assigned a specific IP address that is appropriate to the network to which the computer is attached and that is not assigned to any other computer on that network. If you move a computer to a new network, it must be assigned a new IP address for that new network. You can use DHCP to manage these assignments automatically.

An IP client contacts a DHCP server for configuration parameters. The DHCP server is typically centrally located and operated by the network administrator. Because the server is run by a network administrator, DHCP clients can be reliably and dynamically configured with parameters appropriate to the current network architecture.

You can configure the E Series router to support the following DHCP features:

- DHCP access model
- DHCP proxy client
- DHCP relay

- DHCP relay proxy
- DHCP local server
- DHCP external server

When you enable the DHCP local server, DHCP relay or relay proxy, and DHCP proxy client in the router on the same VR, the feature with the highest order of precedence is always enabled. The order of precedence among these features (highest to lowest) is:

- DHCP local server
- DHCP relay or relay proxy
- DHCP proxy client

Session and Resource Control Software

The Session and Resource Control (SRC) software, formerly the Service Deployment System (SDX) software is a component of Juniper Networks management products. The SRC software provides a Web-based interface that allows subscribers to access services, such as the Internet, an intranet, or an extranet.

When a DHCP subscriber logs in, the SRC software can authorize the address request and select the DHCP address pool on the router from which the DHCP address is selected. The SRC software can also control the number of IP addresses that are given to a particular retailer or subscriber and control the lease time of IP addresses assigned to DHCP subscribers.

The router retrieves the DSL line rate parameters from Access Node Control Protocol (ANCP) and reports this information to the SRC software with the corresponding COPS messages. If the router cannot retrieve the DSL line rate parameters from ANCP, it retrieves the DSL information in the following ways:

- **From AAA layer**—For PPP interfaces, the router retrieves the DSL line rate parameters from the AAA layer and reports this information to the SRC software.
- **From DHCP options**—For DHCP external server and DHCP local server in equal-access mode, the router retrieves the DSL line rate parameters from DHCP options and reports this information to the SRC software. To enable the DHCP external server to receive the DHCP options if the router blocks the DHCP options on the DHCP application, you must use the **set dhcp relay preserve-trusted-client-option** command.



NOTE: The SRC client configured on the E Series router does not send Delete Request (DRQ) messages for interfaces that are bounced during the address mode and are in the administratively up state. Bouncing of an interface refers to shutting down and restarting the interface, releasing the IP address allocated to the clients connected on that interface, and obtaining a fresh IP address for the clients using a rediscovery process. For such interfaces, interface DRQ messages are not sent to the COPS server (or PDP) after DRQ messages for the address configured on the interface are sent from the SRC client.

- Related Documentation**
- *set dhcp relay preserve-trusted-client-option*

DHCP Platform Considerations

For information about modules that support DHCP on the ERX7xx models, ERX14xx models, and the ERX310 Broadband Services Router:

- See *ERX Module Guide, Table 1, ERX Module Combinations* for detailed module specifications.
- See *ERX Module Guide, Appendix A, Module Protocol Support* for information about the modules that support DHCP.

For information about modules that support DHCP on the E120 and E320 Broadband Services Routers:

- See *E120 and E320 Module Guide, Table 1, Module and IOAs* for detailed module specifications.
- See *E120 and E320 Module Guide, Appendix A, IOA Protocol Support* for information about the modules that support DHCP.

- Related Documentation**
- [DHCP Overview Information on page 461](#)
 - [DHCP References on page 463](#)

DHCP References

For more information about DHCP, consult the following resources:

- DSL Forum Technical Report (TR)-101—Migration to Ethernet-Based DSL Aggregation (April 2006)
- RFC 2131—Dynamic Host Configuration Protocol (March 1997)
- RFC 2132—DHCP Options and BOOTP Vendor Extensions (March 1997)
- RFC 3046—DHCP Relay Agent Information Option (January 2001)
- RFC 3315—Dynamic Host Configuration Protocol for IPv6 (DHCPv6) (July 2003)
- RFC 3633—IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) Version 6 (December 2003)
- RFC 4243—Vendor-Specific Information Suboption for the Dynamic Host Configuration Protocol (DHCP) Relay Agent Option (December 2005)



NOTE: IETF drafts are valid for only 6 months from the date of issuance. They must be considered as works in progress. Please refer to the IETF website at <http://www.ietf.org> for the latest drafts.

- Related Documentation**
- [DHCP Overview Information on page 461](#)
 - [DHCP Platform Considerations on page 463](#)
 - [RADIUS IETF Attributes Supported for Subscriber AAA Accounting Messages on page 193](#)
 - [RADIUS IETF Attributes Supported for AAA Tunnel Accounting Messages on page 200](#)
 - [RADIUS IETF Attributes on page 241](#)

Configuring the DHCP Access Model

The E Series router provides a DHCP access model, which enables you to integrate the router into an existing RADIUS-based operation support system (OSS). In the DHCP access model, a DHCP local server or DHCP external service is configured, but the E Series router does not have direct interaction with an OSS or a policy server, such as the Session and Resource Control (SRC) software. The router passes the client's DHCP options, client's media access control (MAC) address and, if appropriate, the DHCP relay's IP address in RADIUS requests for authentication.

To configure the DHCP access model to pass the client's information in RADIUS requests, you enable the DHCP options feature, then specify the client information to be passed to RADIUS. You can specify that the client's MAC address be included in the request. You can also specify that the DHCP relay's IP address be sent, if appropriate. For descriptions of the RADIUS attributes used with the DHCP access model, see [“Juniper Networks VSAs Supported for Subscriber AAA Access Messages” on page 186](#) and [“Juniper Networks VSAs Supported for Subscriber AAA Accounting Messages” on page 196](#).

Configuring DHCP Proxy Clients

Dynamic Host Configuration Protocol (DHCP) proxy client support enables the router to obtain an IP address from a DHCP server for a remote Point-to-Point Protocol (PPP) client. Each virtual router (acting as a DHCP proxy client) can query up to five DHCP servers.

For PPP users, the router acts as a DHCP client to obtain an address for the user. This is referred to as DHCP proxy.

The process for PPP users is as follows:

1. The remote user dials in and the client requests RADIUS authentication.
2. The authentication, authorization, and accounting (AAA) server on the router sends a request to the DHCP proxy client on the router for an IP address to be assigned to the remote user's host.
3. The proxy client assumes the role of a DHCP client and sends a discovery message to each DHCP server.
4. One or more of the DHCP servers responds with an offer message containing an IP address.

5. The proxy client determines which offer to accept and sends a message to that DHCP server requesting that IP address.
6. The DHCP server responds to the proxy client with an acknowledgment message.
7. The proxy client passes the IP address to the AAA server on the router, and the AAA server returns the address to PPP. PPP then assigns the address to the remote host. The new IP address is included when the router next updates its routing table.

Dynamic IP addresses are leased to the remote host for a specific period of time, which can range from minutes to days. At the halfway point in the lease period, the proxy client requests an extension from the DHCP server on behalf of the remote host. The lease is extended for a period specified in the acknowledgment (ACK) message returned by the DHCP server—typically equal to the original lease. If the DHCP server returns a negative acknowledgment (NAK) message to the proxy client, the proxy client notifies the server on the router that the extension has been denied. The AAA server logs out the remote host and frees the IP address for reuse.

When a remote host disconnects, the AAA server notifies the proxy client that the IP address is available for reuse. The proxy client informs the DHCP server, which can now reassign that IP address.



NOTE: The maximum number of DHCP proxy client bindings that are stored on the router chassis is 48,000 for ERX routers, 64,000 for E120 routers, and 96,000 for E320 routers.

For additional information about managing client bindings, see [“Viewing and Deleting DHCP Client Bindings” on page 466](#).

To configure a proxy client from Global Configuration mode:

1. Specify the address of the DHCP server that will provide IP addresses for remote hosts. You can specify a maximum of five DHCP servers.

```
host1(config)#ip dhcp-server 10.6.128.10
```

2. Direct the router to request IP addresses for remote users from the DHCP servers.

```
host1(config)#ip address-pool dhcp
```

Related Documentation

- [ip address-pool](#)
- [ip dhcp-server](#)

Logging DHCP Packet Information

The JunosE Software enables you to collect and log DHCP packet information for all JunosE DHCP access models on a per-interface basis. To log packets for a specific DHCP application, you enable DHCP packet logging on the interface that serves the application. JunosE Software supports per-interface DHCP packet logging on a maximum of 16 interfaces. Per-interface DHCP packet logging is disabled by default.

You can specify which packets are logged—receive, transmit, or all. You can optionally assign low or high priority to the logged packets. Packets are assigned a low priority by default, which does not interfere with router DHCP packet processing. The logged packets are output to the `dhcpCapture` event logging category.

You can configure per-interface DHCP packet logging on statically configured and dynamically created IP interfaces. However, configuration information for dynamic interface configurations is lost after a cold restart. Both static and dynamic interface configuration information is maintained after a warm restart.

You use the **`ip dhcp-capture`** command with the following keywords to enable DHCP packet logging for all DHCP applications on the interface.

- Use the **`receive`**, **`transmit`**, and **`all`** keywords to specify the type of DHCP packets that is logged.
- Use the optional **`priority`** keyword to assign a **`low`** or **`high`** priority to logged packets. By default, logged packets have a low priority and do not interfere with the router's DHCP packet processing.

You can specify DHCP packet logging on a maximum of 16 interfaces.

- To enable DHCP packet logging:

```
host1(config-if)#ip dhcp-capture all
```

**Related
Documentation**

- *ip dhcp-capture*

Viewing and Deleting DHCP Client Bindings

The JunosE Software provides commands that enable you to manage your router's DHCP external server, DHCP local server, and DHCP relay proxy client bindings. A client binding associates an IP address with a DHCP client, and describes both the client (for example, hardware address and state) and the IP address (for example, subnet and lease time).

The following commands enable you to view information about current DHCP client bindings:

- To display information and track lease times and status for specified DHCP client bindings, with results arranged in ascending order by binding ID, use the **`show dhcp binding`** command.
- To display information and track lease times and status for specified DHCP client bindings, with results arranged in ascending order by IP address, use the **`show dhcp host`** command. This command displays information only for DHCP client bindings with assigned IP addresses.
- To display count information for DHCP client bindings and interfaces, use the **`show dhcp count`** command.

To delete a connected user's IP address lease and the associated route configuration when the DHCP client binding is no longer needed, use the **`dhcp delete-binding`** command.

When you delete a DHCP client binding, the lease is removed on the router. You might delete client bindings to:

- Recover functional resources from a user who has not explicitly terminated connectivity and whose lease is unexpired.
- Discontinue connectivity to a user, prompting or forcing the user to request a new lease in order to reestablish network connectivity.

The router does not notify the DHCP client or the DHCP server when you issue the **dhcp delete-binding** command.



NOTE: The **dhcp delete-binding** command replaces the **clear ip dhcp-local binding** and **dhcp-external delete-binding** commands, which are deprecated and might be removed in a future release.

Use the following keywords and variables with the **dhcp delete-binding** command to specify (filter) the client bindings you want to delete:

- **all**—All DHCP local server, DHCP external server, and DHCP relay proxy client bindings
- **all-local**—All DHCP local server client bindings
- **all-external**—All DHCP external server client bindings
- **all-relay-proxy**—All DHCP relay proxy client bindings
- *binding-id*—DHCP binding ID for a specific client
- **circuit-id**—Agent-circuit-id suboption (suboption 1) string of the DHCP relay agent information option (option 82); the circuit ID string supports matching of both regular expression metacharacters and nonprintable ASCII characters in binary sequences
- **external**—DHCP external server bindings that meet the deletion criteria
- **interface**—Interface string associated with DHCP client bindings; the interface string supports matching of regular expression metacharacters, and must be specified as a regular expression without spaces
- **ip-prefix**—IP prefix (address and subnet mask) of the DHCP client
- **local**—DHCP local server bindings that meet the deletion criteria
- **no-interface**—DHCP clients without a lower-layer interface; use this keyword to delete DHCP client bindings configured over dynamic interfaces for which the lower-layer interface has been shut down
- **relay-proxy**—DHCP relay proxy bindings that meet the deletion criteria
- **remote-id**—Agent-remote-id suboption (suboption 2) string of the DHCP relay agent information option (option 82); the remote ID string supports matching of both regular expression metacharacters and nonprintable ASCII characters in binary sequences
- *subnetAddress*—IP address of the subnet on which the DHCP client resides

Filtering the deletion of DHCP client bindings by the circuit ID string or remote ID string is not supported for the DHCP external server application. DHCP external server does not store information about the agent-circuit-id suboption or agent-remote-id suboption of option 82.

You can remove all DHCP client bindings, all DHCP client bindings of a particular type, or a specified DHCP client binding that meets the deletion criteria you specify.

- To delete all DHCP client bindings on virtual router vr1:

```
host1:vr1#dhcp delete-binding all
```

- To delete DHCP local server client bindings with the specified subnet address:

```
host1:vr2#dhcp delete-binding local 0.0.0.0
```

When you delete DHCP client bindings of a particular type on a specified subnet, you must specify the **local**, **external**, or **relay-proxy** type keyword to prevent accidental deletion of all DHCP client bindings.

- To delete a specific DHCP client binding:

```
host1:vr1#dhcp delete-binding 3972819365
```

- To delete DHCP client bindings with the specified IP prefix:

```
host1:vr1#dhcp delete-binding ip-prefix 10.1.0.0/28
```

- To delete DHCP client bindings without a lower-layer interface:

```
host1:vr1#dhcp delete-binding no-interface
```

- To delete DHCP client bindings with the specified interface string:

```
host1:vr2#dhcp delete-binding interface ip71.*4
```

This **dhcp delete-binding** command uses the * (asterisk) regular expression metacharacter in the interface string to delete DHCP client bindings on virtual router vr2 with an IP address beginning with 71 and ending with 4.

- To delete DHCP client bindings that match the specified circuit ID string:

```
host1:vr3#dhcp delete-binding circuit-id \xe3
```

To specify nonprintable byte codes in the circuit ID string or remote ID string, you can use the string `\xab`, where `ab` is a hex code of the byte. This **dhcp delete-binding** command uses the string `\xe3` to represent byte E3 in the circuit ID string. This command deletes DHCP client bindings on virtual router vr3 with the specified circuit ID string.

**Related
Documentation**

- *dhcp delete-binding*
- *show dhcp binding*
- *show dhcp count*
- *show dhcp host*

DHCP Client Bindings and Duplicate MAC Addresses for Subinterfaces Overview

In certain network scenarios, active VLAN subinterfaces of subscribers might be transferred from one virtual router to another, and later retransitioned to the original virtual router for correct computation of subscription and billing costs for customers being serviced by an enterprise provider. Also, addition and removal of active VLAN subinterfaces might be performed during troubleshooting with the customer premises equipment (CPE) devices. Such changes in the configuration of active VLAN subinterfaces causes differences in the subscriber entries displayed in the output of the **show dhcp bindings** (and other commands used to monitor DHCP bindings) and **show subscribers** commands.

When the DHCP client is bound to an IP address, deletion of the active VLAN subinterface causes the subscriber entry to be removed from the AAA database and the access-internal route for that client to be deleted. In such a scenario, if the client binding was still retained in the DHCP database, the entries for that subscriber for which the binding is removed from the AAA database are not displayed in the output of the **show subscribers** (under the Username field) and **show ip route access-internal** (under the Prefix/Length field) commands.

When the VLAN subinterface associated with a DHCP client, which was previously deleted when the client binding was removed, is reconfigured, the entries for that subscriber are not displayed in the output of the following **show** commands until the DHCP client sends a discover or renew request to the DHCP server for an IP address to be allocated to it:

- **show ip dhcp-local binding interface** (under the Address field)
- **show ip route access-internal** (under the Prefix/Length field)
- **show subscribers** (under the Username field)

When some DHCP packets flow between the subscriber and the router, the following events take place:

- During the process of allocating IP addresses to the DHCP client, which involves the discovery, offer, request, and acknowledgment messages between the server and the client, the client binding already exists in the database and the DHCP server does not contact AAA for authentication. At this point, the subscriber entry is not present in the AAA database. The access-internal route is created for the client and the subscriber connection becomes active. The client does not receive Acct-Request packets because the entry for this subscriber is not available in the AAA database.
- When the client sends a renew request to renew its address, the request does not reach the interface on the DHCP server. The DHCP server sends a NAK message to the client, forcing the client to begin the DHCP connection process again.
- When the client sends a rebind request for the IP address to be bound again to it, the existing binding for this client is deleted and re-created during the next discovery process. All the databases are synchronized and the entry for the client is correctly displayed in the output of the **show subscribers** and **show dhcp bindings** commands.

In this scenario, the subscriber session might be established and active without accounting records for Acct-Stop and Interim-Acct messages sent to the RADIUS server during the

process of allocating addresses to DHCP clients in JunosE releases numbered lower than 9.3.x.

Beginning with JunosE Release 9.3.x, support for configuring DHCP external server to uniquely identify clients with duplicate MAC addresses is available. This functionality causes a new IP address to be assigned to a client during the process of DHCP address allocation by the DHCP server using the discovery, offer, request, and acknowledgment sequence. The previously configured binding for the same client is deleted from the database before the lease period expires for that address, immediately after the VLAN subinterface for that client is deleted. Because the DHCP bindings are stored in a server management table that includes the VLAN subinterface user ID (UID), when the server queries the management table to check whether a binding for a client already exists, no match is found and a fresh client binding is created when the VLAN subinterface is reconfigured.

To prevent the problem of incorrect and inconsistent parameters being displayed in the **show** commands used to monitor subscriber information and DHCP binding attributes, the client binding is removed from the DHCP database after the VLAN subinterface associated with that subscriber is deleted. Retaining the client binding is not effective after the primary interface is deleted because when the client logs in again, it is assigned a different user ID unless a rollover of the user ID occurs. This rollover causes the user ID assigned to the client prior to the logout to be reassigned to it upon logging in again and a fresh IP address is bound to the client. When a stateful SRP switchover operation is performed before the transaction is posted to the standby SRP module, the client binding remains in the database because it is added again when the configuration data is restored from the mirrored containers. The client binding stays in the database until its lease expires.

- Related Documentation**
- [DHCP Platform Considerations on page 463](#)
 - [DHCP References on page 463](#)

CHAPTER 19

DHCP Local Server Overview

This chapter provides an overview of the DHCP local server on the E Series router. This chapter contains the following sections:

- [Embedded DHCP Local Server Overview on page 471](#)
- [Equal-Access Mode Overview on page 472](#)
- [Standalone Mode Overview on page 474](#)
- [DHCP Local Server Prerequisites on page 477](#)
- [DHCP Local Server Configuration Tasks on page 477](#)
- [DHCP Unique ID for Clients and Servers Overview on page 478](#)
- [Authentication and Accounting of IPv6 Subscribers Using the DHCPv6 Local Server Overview on page 480](#)
- [Interoperation of Authentication of IPv6 Clients and Display of Active Subscriber Information on page 482](#)
- [Limiting the Maximum Number of IPv6 Prefixes Delegated Per Interface by the DHCPv6 Local Server Overview on page 483](#)

Embedded DHCP Local Server Overview

The router offers an embedded DHCP server, known as the DHCP local server. The DHCP local server has two modes: equal-access and standalone.



NOTE: E Series routers also support an embedded DHCP version 6 (DHCPv6) local server. The DHCPv6 local server provides a subset of the features of the DHCP local server. For information about configuring the DHCPv6 local server, see [“Configuring the DHCPv6 Local Server” on page 501](#).

- In equal-access mode, the DHCP local server works with the Juniper Networks SRC software to provide an advanced subscriber configuration and management service.
- In standalone mode, the DHCP local server provides a basic DHCP service, and also allows you to configure AAA authentication for incoming DHCP clients. Also, after successful authentication, the DHCP local server uses the information in the client's AAA subscriber record together with the client's DHCP parameters to select the IP address pool used for address assignment.

DHCP local server also supports RADIUS accounting, including interim accounting, in standalone mode. This feature allows you to use RADIUS start and stop attributes to track user events such as the lifetime of an IP address.

DHCP Local Server and Client Configuration

You can use DHCP to configure the router to allow remote access to non-PPP clients. DHCP-based access is also an alternative to PPP in environments such as Public Wireless LANs (PWLANS). In PWLANS, a user scans for available broadband networks, then is redirected to a web-based authentication mechanism to request service.

DHCP provides address assignment information for users. Authentication, authorization, and accounting are separate processes, and are up to the Internet service provider (ISP) to define.

The DHCP local server can configure a client with the following DHCP options:

- Default router
- DNS server
- Domain name
- Lease time
- Grace period for address lease
- NetBIOS name server
- NetBIOS node type
- Subnet mask

Related Documentation

- [Equal-Access Mode Overview on page 472](#)
- [Standalone Mode Overview on page 474](#)
- [DHCP Local Server Prerequisites on page 477](#)
- [DHCP Local Server Configuration Tasks on page 477](#)
- [Viewing and Deleting DHCP Client Bindings on page 466](#)

Equal-Access Mode Overview

In equal-access mode, the router enables access to non-PPP users. Non-PPP equal access requires the use of the router's DHCP local server and SRC software, which communicates with a RADIUS server.

The DHCP local server performs the following functions in equal-access mode:

- Communicates with SRC software.
- Assigns an IP address that enables the subscriber to access services.

Local Pool Selection and Address Allocation

The DHCP local server selects a DHCP pool from which to allocate an address using a number of parameters in a certain predefined sequence. The router compares the parameters with the local DHCP pools in the order presented in [Table 107 on page 473](#). When the router finds a match, it selects a pool based on the match and does not examine other parameters.

Table 107: Local Pool Selection in Equal-Access Mode

Field	How the DHCP Local Server Uses the Field
Framed IP address	<p>The client's entry can be configured with a framed IP address, which the DHCP local server can get from the SRC software (formerly the SDX software).</p> <p>If the router selects a pool using a framed IP address, the DHCP local server attempts to allocate the framed IP address from the pool. If the framed IP address is not available, then the server allocates the next available address in the pool to the client.</p>
Pool name	Each DHCP local pool has a pool name. The client's entry can also be configured with a pool name, which the DHCP local server can get from the SRC software. The SRC software must be configured to send RADIUS attributes to DHCP.
Domain name	<p>You can use a domain name as the name of a DHCP local pool. If the client logs onto the SRC software and RADIUS authenticates the client using a domain name, the DHCP local server receives the domain name from the SRC software.</p> <p>If the client's domain name does not match the name of the DHCP local pool, the router attempts to match the client's domain name to the domain name field within the pool.</p>
Giaddr	A DHCP local pool is configured with a network address. A gateway IP address (giaddr), which indicates a client's subnet, can be presented to the DHCP local server in the client's DHCP request message. The giaddr field in the DHCP request message contains the IP address of a DHCP relay agent. The router attempts to match the giaddr address in the DHCP request message with the network address of a DHCP local pool.
Received interface IP address	The router uses the IP address of the interface on which the DHCP packet is being processed and attempts to match it with the network address of a DHCP local pool. If the interface address matches with the IP address configured in the DHCP local address pool on the router, that pool is used to delegate the address to the client.

The Connection Process

The following sequence describes how the subscriber connects to the network for the first time using equal-access mode. [Figure 11 on page 474](#) illustrates the process.

1. The subscriber's computer boots and issues a DHCP request.
2. The DHCP local server uses the SRC client to issue a COPS request to retrieve address pool information.
3. After standard DHCP negotiations, the DHCP local server supplies an IP address to the subscriber's computer from a local address pool, as described in the previous section.


The router maintains a host route that maps the IP address to the router's interface associated with the subscriber's computer.

4. The subscriber's computer retains the IP address until the subscriber turns off the computer.



NOTE: If a DHCP client attempts to renew its address and the DHCP server receives the request on a different interface than the interface that the client originally used, the DHCP server sends a NAK message to the client, forcing the client to begin the DHCP connection process again.

Figure 11: Non-PPP Equal Access via the Router



- Related Documentation**
- [Embedded DHCP Local Server Overview on page 471](#)
 - [Standalone Mode Overview on page 474](#)
 - [DHCP Local Server Prerequisites on page 477](#)
 - [DHCP Local Server Configuration Tasks on page 477](#)

Standalone Mode Overview

In standalone mode, the DHCP local server operates as a basic DHCP server. Clients are not authenticated by default; however, you can optionally configure the DHCP local server to use AAA authentication for the incoming clients. The DHCP local server receives DHCP client requests for addresses, selects DHCP local pools from which to allocate addresses, distributes addresses to the clients, and maintains the resulting DHCP bindings in a server management table.

Local Pool Selection and Address Allocation

In standalone mode, the DHCP local server selects a pool to allocate an address for a client; the SRC software is never notified or queried. The process used depends on whether AAA authentication is configured.

- If AAA authentication is not configured, the DHCP local server selects a pool by matching the local pool network address to the giaddr or the received interface IP address. The router compares the parameters with the local DHCP pools in the order presented in [Table 108 on page 475](#). When the router finds a match, it selects a pool based on the match and does not examine other parameters.

Table 108: Local Pool Selection in Standalone Mode Without AAA Authentication

Field	How the DHCP Local Server Uses the Field
Giaddr	A giaddr, which indicates a client's subnetwork, can be presented to the DHCP local server in the client DHCP REQUEST message. The giaddr field in the DHCP request message usually contains the IP address of a DHCP relay agent. The router attempts to match the giaddr address in the DHCP request message with the network address of a DHCP local pool. If it finds a match, the router uses the matching DHCP local pool.
Received interface IP address	The router uses the IP address of the interface on which the DHCP packet is being processed and attempts to match it with the network address of a DHCP local pool.

After the router selects a DHCP local pool, the DHCP local server first tries to find a reserved IP address for the client in the selected pool. If no reserved address is available, the router attempts to allocate a client's requested IP address. If the requested IP address is not available, the router allocates the next available address in the pool. If a grace period is configured for the pool, the router assigns the grace period to the allocated address. If no addresses are available in a pool, the DHCP local server attempts to allocate an address from the linked pool, if such a pool is configured.

- If AAA authentication is configured (as described in ["Configuring AAA Authentication for DHCP Local Server Standalone Mode" on page 497](#)) and the authentication is successful, the local server selects an IP address pool based on the order presented in [Table 109 on page 475](#). When the router finds a match, it selects a pool based on the match and does not examine other parameters.

Table 109: Local Pool Selection in Standalone Mode with AAA Authentication

Field	How the DHCP Local Server Uses the Field
Framed IP address	<p>The client's RADIUS entry can be configured with a framed IP address, which the DHCP local server can get from the AAA server when the client is authenticated.</p> <p>If the AAA server specifies a framed IP address, the DHCP local server attempts to allocate the address pool that contains the framed IP address and allocates that address from the pool.</p>

Table 109: Local Pool Selection in Standalone Mode with AAA Authentication (*continued*)

Field	How the DHCP Local Server Uses the Field
Pool name	<p>Each DHCP local pool has a pool name. The client's RADIUS entry can also be configured with a pool name, which the DHCP local server can get from the AAA server when the client is authenticated. The AAA server must be configured to send RADIUS attributes to DHCP.</p> <p>If AAA specifies an address pool name, the local server finds the pool with the matching name and allocates an address from that pool.</p>
Domain name	<p>You can use a domain name as the name of a DHCP local pool. If RADIUS authenticates the client using a domain name, the DHCP local server receives the domain name from the AAA server.</p> <p>If the client's domain name does not match the name of the DHCP local pool, the router attempts to match the client's domain name to the domain name field within the pool.</p>
Giaddr	A DHCP local pool is configured with a network address. A gateway IP address (giaddr), which indicates a client's subnetwork, can be presented to the DHCP local server in the client's DHCP request message. The giaddr field in the DHCP request message usually contains the IP address of a DHCP relay server. The router attempts to match the giaddr address in the DHCP request message with the network address of a DHCP local pool.
Received interface IP address	The router uses the IP address of the interface on which the DHCP packet is being processed and attempts to match it with the network address of a DHCP local pool. If the interface address matches with the IP address configured in the DHCP local address pool on the router, that pool is used to delegate the address to the client.

After the router selects a DHCP local pool, the DHCP local server performs the following functions for allocating an address to the client:

- If the DHCP local pool is selected using the framed IP address, the server attempts to allocate the framed IP address from the pool. If the framed IP address is in use, the server attempts to allocate an address from the linked pool (if configured).
- If the DHCP local pool is selected using fields (such as pool name) other than the framed IP address, the server first tries to find the reserved IP address for the client in the selected pool. If the reserved address is in use, the server attempts to allocate the IP address requested by the client. If the requested IP address is in use, the server allocates the next available address in the selected pool. If all addresses in the selected pool are in use, the server attempts to allocate an address from the linked pool (if configured).

Server Management Table

For each client that makes requests of the DHCP local server, the router keeps an entry in the server management table. The entry defines client-specific information and state information. The router uses this table to identify clients when it receives subsequent messages and to maintain the state of each client within the DHCP protocol. In addition, the table contains information that may be transferred to and from the SRC software.

Related Documentation

- [Embedded DHCP Local Server Overview on page 471](#)

- [Equal-Access Mode Overview on page 472](#)
- [DHCP Local Server Prerequisites on page 477](#)
- [DHCP Local Server Configuration Tasks on page 477](#)

DHCP Local Server Prerequisites

Before you configure DHCP local server, you need to configure interfaces. You can configure ATM or Ethernet interfaces for DHCP local server. These interfaces can be numbered or unnumbered. Because subscribers connect to the router from different subnetworks, you must configure an IP address for each subnetwork on the interface. This action provides connectivity between the subnetwork and the router.

To configure a numbered IP address for DHCP local server:

1. Select an ATM or Ethernet interface.
2. Assign the primary IP address for one subnetwork to this interface.
3. Assign secondary IP addresses for all other subnetworks to this interface.

To configure an unnumbered IP address for DHCP local server:

1. Specify a loopback interface.
2. Assign the primary IP address for one subnet to the loopback interface.
3. Assign secondary IP addresses for all other subnets to the loopback interface.
4. Select an ATM or Ethernet interface.
5. Configure an unnumbered IP address associated with the loopback interface on the ATM or Ethernet interface.

For information about defining IP addresses, see the *Configuring IP* chapter in the *JunosE IP, IPv6, and IGP Configuration Guide*.

- Related Documentation**
- [IP Profiles](#)
 - [Setting Up an Unnumbered Interface](#)

DHCP Local Server Configuration Tasks

This section covers the configuration tasks for equal-access and standalone modes. Perform the appropriate procedure:

1. For both equal-access and standalone modes, configure the DHCP local server.
[See “Configuring AAA Authentication for DHCP Local Server Standalone Mode” on page 497](#) for a sample configuration.
2. For standalone mode, optionally configure the router to use AAA authentication for DHCP requests from subscribers.

See [“Configuring AAA Authentication for DHCP Local Server Standalone Mode”](#) on page 497 for a sample configuration.

3. For non-PPP equal access, configure the router to work with the SRC software.

See [“Configuring the Router to Work with the SRC Software”](#) on page 507 for a sample configuration.

**Related
Documentation**

- [Embedded DHCP Local Server Overview on page 471](#)
- [Equal-Access Mode Overview on page 472](#)
- [Standalone Mode Overview on page 474](#)

DHCP Unique ID for Clients and Servers Overview

Each entity in a DHCP operation, the client and the server, has a DHCP unique identifier (DUID). DHCP clients use DUIDs to identify a server in messages where a server needs to be identified. DHCP servers use DUIDs to determine the configuration parameters to be used for clients and in the association of addresses with clients.

The DUID is contained in the client identifier and server identifier options. The DUID is stable for any specific client or server. The DHCPv6 application uses DUIDs based on link-layer addresses for both the client and server identifier options. An Identity Association for Prefix Delegation option is a collection of prefixes assigned to a requesting router. A requesting router can have more than one Identity Association for Prefix Delegation option; for example, one for each of its interfaces. Each Identity Association for Prefix Delegation is denoted by an Identity Association identifier. The Identity Association identifier is chosen by the requesting router and is unique among the Identity Association identifiers that are present in the Identity Association for Prefix Delegation options on the requesting router. A client binding is indexed by a DUID.

When an IPv6 DHCP client requests two prefixes with the same DUID but different Identity Association identifiers on two different interfaces, these prefixes are considered to be for two different clients, and the interface information is maintained for both the clients. Clients and servers identify DUIDs as opaque values and compare DUIDs only to check for their equality. Clients and servers do not process DUIDs for other information.

A DUID consists of a two-octet type code represented in network byte order, followed by a variable number of octets that make up the actual identifier; for example, 00:02:00:01:02:03:04:05:07:a0. A DUID can be up to 128 octets in length (excluding the Type code). The following types are currently defined for the DUID parameter:

- Type 1—Link-layer address plus time (DUID-LLT)
- Type 2—Vendor-assigned unique ID based on Enterprise Number (DUID-EN)
- Type 3—Link-layer address (DUID-LL)

The Type 1 DUID consists of a two-octet type field that contains the value 1, a two-octet hardware type code, four octets that signify a time value, followed by the link-layer

address of any one network interface that is connected to the DHCP device at the time that the DUID is generated.

The Type 2 DUID is assigned by the vendor to the device and contains the vendor's registered private enterprise number as maintained by the IANA, followed by a unique identifier assigned by the vendor.

The Type 3 DUID contains a two-octet type field that stores the value 3, a two-octet network hardware type code, followed by the link-layer address of any one network interface that is permanently connected to the client or server device.

By default, the DHCPv6 local server application in JunosE Software uses the Type 2 server DUID for allocation of IPv6 prefixes from the delegating router, which is an E Series router configured as a DHCPv6 local server to requesting routers, which is the customer premises equipment (CPE) at the edge of the remote client site that acts as the DHCP client. In scenarios in which the CPE does not support the Type 2 DUID, or if the service provider uses a DUID type other than Type 2, the verification of identity of servers and clients by each other using DUIDs does not happen successfully. In such network environments, configuring the DUID type on the DHCPv6 local server to be other than the default value of Type 2 enables correct identity verification of clients and servers.

You can configure the type of DUID using the **ipv6 dhcpv6-local duid-type** *duidType* command in Global Configuration mode to be either Type 2 or Type 3. These two types are currently supported by the DHCPv6 local server application in JunosE Software. The Type 1 DUID is not supported by the DHCPv6 local server in JunosE Software. However, DHCPv6 clients support DUID Types 1, 2, and 3. The **ipv6 dhcpv6-local duid-type** command enables you to specify the DUID type that matches with the DUID type that the service providers use in their networks and also provides flexibility to DHCP subscribers to use a DUID type that suits their requirements. The DHCPv6 local server uses the configured DUID type in its communication with the client.

The DUID type conforms to the following guidelines:

- The DUID type is unique across all the virtual routers on the chassis.
- The DUID type is persistent across a system reload.
- The DUID type is retained after a switchover.
- You cannot modify the DUID type when at least one active DHCP client connection exists.

To support the Type 3 DUID, the DHCPv6 local server uses a combination of the chassis ID and virtual router ID as the DUID. When you remove the configured DUID type using the **no ipv6 dhcpv6-local duid-type** command, the router reverts to the default Type 2 DUID. All the binding requests from the clients are acknowledged with the default server ID if the Type 2 DUID is on the DHCPv6 local server.

Related Documentation

- [Configuring the Type of DHCP Unique ID for DHCPv6 Local Servers on page 504](#)
- *ipv6 dhcpv6-local duid-type*

Authentication and Accounting of IPv6 Subscribers Using the DHCPv6 Local Server Overview

You can use the DHCPv6 local server to perform authentication and accounting of IPv6 subscribers that are directly connected using Ethernet VLAN links to the router. For PPP subscribers, authentication and accounting operations are performed by the underlying PPP module; the DHCPv6 local server only delegates IPv6 prefixes to requesting clients. IPv6 subscribers that are connected over PPP links and IPv6 subscribers that are connected over Ethernet and VLAN interfaces can coexist on a virtual router when you configure the DHCPv6 local server for standalone mode with AAA authentication.

For PPP subscribers, the PPP module authenticates users during the establishment of the PPP session and sends the authentication token to the DHCPv6 local server for allocation of IPv6 prefixes. For IPv6 subscribers, the DHCPv6 local server performs the AAA authentication of clients that are logging in. Prefix delegation for IPv6 subscribers occurs only if the prefix is configured on the interface or if the interface address matches with any of the prefix ranges configured in the IPv6 local address pool on the router. When you configure standalone mode with AAA authentication for the DHCPv6 local server, delegation of prefixes is performed based on the Access-Accept and Access-Reject messages the AAA server sends in response to the client authentication request.

The DHCPv6 local server enables you to optionally configure AAA-based authentication of standalone mode DHCPv6 clients. By default, clients are not authenticated in standalone mode. Typically, an incoming DHCPv6 client does not provide a username—therefore, the DHCPv6 local server constructs a username based on the user's attachment parameters and optional DHCP parameters. AAA uses the constructed username to authenticate the incoming client and create the AAA subscriber record for the client. The information in the AAA subscriber record is then used to determine the IP address pool from which to assign the address for the DHCP client.

You can include the following parameters in the username:

- User prefix
- Circuit type
- Circuit identifier
- Domain name

The complete format of the username is as follows:

user-prefix.circuit-type.circuit-identifier@domain

The elements of the username are defined as follows:

- user-prefix—A configured string per DHCPv6 local server.
- circuit-type—Specifies the circuit type of the interface on which the DHCPv6 client's request was received. Possible values are atm, eth, or vlan.
- circuit-identifier—Specifies the circuit identifier of the interface on which the DHCPv6 client's request was received. The interface identifier has one of the following formats:

- atm—slot.port.vpi.vci
- eth—slot.port.0.0
- vlan—slot.port.svlan.vlan
- domain—Name of the user domain for each DHCPv6 local server.

You can construct the username by using only the user-prefix attribute, using a combination of the user-prefix and domain attributes, or using other optional attributes that are specified. If you remove the domain configuration, the '@' character is removed from the username. The username is valid only when the nondomain portion consists of at least one character, either using the configuration of a non-null user-prefix or using the inclusion of at least one optional username parameter.

The authentication process starts before the Advertise message is sent from the DHCPv6 local server to the client. If the authentication of the subscriber is successful, the DHCPv6 local server sends the Advertise packet to the client in response to DHCPv6 Solicit messages that are received from the client. When the authentication request is sent to the AAA server, the DHCPv6 local server includes the constructed username, password, interface ID, authentication type, and the interface on which the request was received from the user. The AAA server uses this information during authentication and accounting updates. The authentication and accounting attributes that are sent to the RADIUS server are based on RADIUS attributes configured for inclusion in RADIUS messages using the **radius include** command.

This mode of operation for the DHCPv6 local server is called standalone mode with AAA authentication. The default mode of operation for the DHCPv6 local server is standalone mode without AAA authentication that interoperates with the existing capabilities of PPP and non-PPP subscribers.

Accounting for IPv6 Subscribers with DHCPv6 Local Server Standalone Mode

The PPP application handles the transmission of accounting information to the AAA server. The DHCPv6 local server uses the authentication token that the AAA server generates while authenticating the IPv6 user to send the interim accounting updates to the AAA server. The starting and termination of accounting is performed during the authentication phase based on the receipt of the DHCP release packets from clients or the lease expiration of the assigned address.

The Acct-Start message is sent to the RADIUS server after the AAA server receives the message about successful authentication. You can use the **aaa service accounting interval** command to specify the default service interim accounting interval. Service Manager uses this interval value for service accounting when the Service-Interim-Acct-Interval attribute is not configured. Based on the configured interval, the DHCPv6 local server generates interim accounting information. The Acct-Stop message is sent to the RADIUS server when a client binding is removed. The accounting functionality of the DHCPv6 local server is similar to the accounting operations of the DHCPv4 local server.

Related Documentation

- [Interoperation of Authentication of IPv6 Clients and Display of Active Subscriber Information on page 482](#)

- [Configuring AAA Authentication for DHCPv6 Local Server Standalone Mode on page 500](#)
- [Monitoring DHCPv6 Local Server Authentication Information on page 596](#)

Interoperation of Authentication of IPv6 Clients and Display of Active Subscriber Information

The following cases describe the behavior of the **show subscribers** command, used to view details on active subscribers logged in to the router, when AAA and RADIUS authentication mechanisms are used to authenticate IPv6 subscribers:

- If you do not disable AAA authentication on the default router for IP subscribers by using the **aaa authentication ip default none** command and do not map the domain name of the user with the virtual router by using the **aaa domain-map** command, details regarding the logged-in subscribers are not displayed in the output of the **show subscribers** and **show subscribers ipv6** commands. In such cases, you can use the output of the **show subscribers summary** command to view the summary information of the subscribers that are logged in.
- If you disable AAA authentication on the default router for IP subscribers by entering the **aaa authentication ip default none** command and map the domain name of the user with the virtual router by using the **aaa domain-map** command, you can view the details of the active subscribers using the output of the **show subscribers**, **show subscribers ipv6**, and **show subscribers summary** commands.
- If you configure RADIUS authentication for IP subscribers on the default router by using the **aaa authentication ip default radius** command and the Virtual-Router VSA attribute [26-1] is returned from the RADIUS server in the Access-Accept message, the subscriber details are displayed in the output of the **show subscribers** command. Also, you can use the **show subscribers summary** command to view the consolidated information on active subscribers.
- If you configure RADIUS authentication for IP subscribers on the default router by using the **aaa authentication ip default radius** command and the Ipv6-Virtual-Router VSA attribute [26-45] is returned from the RADIUS server in the Access-Accept message, the subscriber details are displayed in the output of the **show subscribers ipv6** command. Also, you can use the **show subscribers summary** command to view the consolidated information on active subscribers.

Related Documentation

- [Authentication and Accounting of IPv6 Subscribers Using the DHCPv6 Local Server Overview on page 480](#)
- [Configuring AAA Authentication for DHCPv6 Local Server Standalone Mode on page 500](#)
- [Monitoring DHCPv6 Local Server Authentication Information on page 596](#)
- *show subscribers*
- *aaa authentication default*
- *aaa domain-map*

Limiting the Maximum Number of IPv6 Prefixes Delegated Per Interface by the DHCPv6 Local Server Overview

You can limit the maximum number of IP version 6 (IPv6) prefixes offered per interface by the Dynamic Host Configuration Protocol version 6 (DHCPv6) local server to DHCPv6 subscribers. The limit value for the offered IPv6 prefixes can be set as a global limit for all interfaces of a particular interface type or as an individual limit for a specific interface or subinterface. You can use the **ipv6 dhcpv6-local limit** command to set the limit for the number of IPv6 prefixes offered by the DHCPv6 local server. With the **ipv6 dhcpv6-local limit** command, you can set the global limit for the following interface types:

- Asynchronous Transfer Mode (ATM)
- Ethernet
- Virtual LAN (VLAN)
- Packet over SONET



NOTE: The DHCPv6 local server is not supported over the packet over SONET (POS) access interface.

When you set the limit value as a global limit for an interface type, all dynamically or statically created interfaces of that type inherit the global limit value. By default, the limit value is set as 48,000 for ATM, Ethernet, and VLAN interface types.

You can override the inherited global limit value for a particular interface by configuring an individual limit for that interface. You can use the **interface** keyword with the **ipv6 dhcpv6-local limit** command to configure the individual limit.

The global limits for interface types and the individual limits set on static interfaces are stored in nonvolatile storage (NVS). These values are restored during a switchover or a reload.

When you assign an individual limit to a dynamic interface, that limit is in force only until either a switchover or reload takes place. After the switchover or reload, if the action that caused the dynamic interface to be created occurs again, a new dynamic interface is created. The new dynamic interface then inherits the limit set by the global values based on the type of interface that is created.



NOTE: Limits that you specify on dynamic interfaces are not restored after a switchover or reboot.

A subinterface table is maintained for both DHCPv6 standalone and authenticate modes to validate the IPv6 prefix limit per interface when the subscriber tries to connect with the local server. The table contains the following information:

- Number of subscribers connected
- Number of subscribers denied upon exceeding the configured prefix limit
- Total number of subscribers attempted
- Type of the subinterface

**NOTE:**

- When you upgrade the JunosE Software by using the unified ISSU from the old release not containing this feature to the new release containing this feature, the subscriber count is retrieved and added to the subinterface table.
 - The IPv6 prefix limit per interface is supported only for the existing Identity Association Prefix Delegation option.
-

**Related
Documentation**

- [Configuring the DHCPv6 Local Server on page 501](#)
- *ipv6 dhcpv6-local limit*

CHAPTER 20

Configuring DHCP Local Server

This chapter provides information for configuring the DHCP local server on the E Series Broadband Services Routers. This chapter contains the following sections:

- [Configuring the DHCP Local Server on page 485](#)
- [Configuring DHCP Local Address Pools on page 491](#)
- [Sending RADIUS Attributes to DHCP Subscribers Overview on page 495](#)
- [Configuring AAA Authentication for DHCP Local Server Standalone Mode on page 497](#)
- [Configuring AAA Authentication for DHCPv6 Local Server Standalone Mode on page 500](#)
- [Configuring the DHCPv6 Local Server on page 501](#)
- [Configuring the Type of DHCP Unique ID for DHCPv6 Local Servers on page 504](#)
- [Deleting DHCPv6 Client Bindings on page 505](#)
- [Configuring the Router to Work with the SRC Software on page 507](#)

Configuring the DHCP Local Server

Tasks to configure the DHCP local server include:

- [Basic Configuration of DHCP Local Server on page 485](#)
- [Limiting the Number of IP Addresses Supplied by DHCP Local Server on page 487](#)
- [Excluding IP Addresses from Address Pools on page 487](#)
- [Configuring DHCP Local Server to Support Creation of Dynamic Subscriber Interfaces on page 488](#)
- [Differentiating Between Clients with the Same Client ID or Hardware Address on page 488](#)
- [Logging Out DHCP Local Server Subscribers on page 489](#)
- [Clearing an IP DHCP Local Server Binding on page 490](#)
- [Using SNMP Traps to Monitor DHCP Local Server Events on page 490](#)
- [Using DHCP Local Server Event Logs on page 491](#)

Basic Configuration of DHCP Local Server

Before you configure a DHCP local server, you must identify which mode to activate (standalone mode or equal-access mode). Use equal-access mode, if you use the Session and Resource Control (SRC) software for address allocation and managing the

subscribers. If you do not use SRC for managing subscribers, use standalone mode. SRC contributes to the address pool selection and so when you use standalone mode, SRC is not used for address allocation.

If you do not specify a mode, equal-access mode is activated, by default. When you activate equal-access mode, common open policy service usage for policy provisioning (COPS-PR) and SRC client are automatically started on the virtual router.

To configure the DHCP local server:

1. Enable the DHCP local server for either equal-access or standalone mode.

```
host1(config)#service dhcp-local equal-access
host1(config)#service dhcp-local standalone
```

2. (Optional) Specify the maximum number of IP addresses that the DHCP local server can supply to each VPI/VCI, VLAN, Ethernet subnetwork, or to a particular interface or subinterface. See [“Limiting the Number of IP Addresses Supplied by DHCP Local Server” on page 487](#) for more information about limiting the number of IP addresses.

```
host1(config)#ip dhcp-local limit ethernet 6
```

3. (Optional) Specify any addresses that the DHCP local server must not assign. See [“Excluding IP Addresses from Address Pools” on page 487](#) for more information.



NOTE: You can specify this command multiple times on the CLI and the excluded address must fall within a network that has been specified in the DHCP local pool.

```
host1(config)#ip dhcp-local excluded-address 10.10.3.4
```

4. (Optional) Enable general DHCP local server traps. See [“Using SNMP Traps to Monitor DHCP Local Server Events” on page 490](#).

```
host1(config)#ip dhcp-local snmpTraps
```

5. (Optional) Configure the DHCP local server to support the creation of dynamic subscriber interfaces built over dynamic VLANs that are based on the agent-circuit-id option (suboption 1) of the option 82 field in DHCP messages. See [“Configuring DHCP Local Server to Support Creation of Dynamic Subscriber Interfaces” on page 488](#).

```
host1(config)#ip dhcp-local auto-configure agent-circuit-identifier
```

6. (Optional) Specify that DHCP local server use an optional method to differentiate between clients with duplicate client IDs or hardware addresses. Any changes you make have no effect on currently bound clients. See [“Differentiating Between Clients with the Same Client ID or Hardware Address” on page 488](#).

```
host1(config)# ip dhcp-local unique-client-ids
```

7. Configure the DHCP local address pool that supplies IP addresses to subscribers who want to access a domain. See [“Configuring DHCP Local Address Pools” on page 491](#) for more information about configuring address pools.

Limiting the Number of IP Addresses Supplied by DHCP Local Server

You can specify the maximum number of IP addresses that the DHCP local server can supply to each VPI/VCI, VLAN, Ethernet subnetwork, or POS access interface type, or to a particular interface or subinterface.

You can set global limits for a given interface type—all interfaces of that type that are subsequently created, whether dynamically or statically, inherit that limit value.

You can also set an individual interface limit for a specific interface and override the global limit configured for that interface type. For example, suppose the VLAN interface type limit is five. You can specify a limit of 10 for the VLAN interface FastEthernet 1/0.100. All other VLAN interfaces retain the global limit of five.

The global limits for interface types and the individual interface limits set on static interfaces are kept in NVS. These values are restored during a switchover or a reload.

When you assign an individual limit to a dynamic interface, that limit is in force only until either a switchover or reload takes place. After the switchover or reload, if the action that caused the dynamic interface to be created occurs again, a new dynamic interface is created. The new dynamic interface then inherits the limit set by the global values based on the type of interface that is created.

- To set a global limit for an interface type:

```
host1(config)#ip dhcp-local limit ethernet 6
```

- To set a limit for a specific interface:

```
host1(config)#ip dhcp-local limit interface atm 3/1 15
```



NOTE: Limits that you specify on dynamic interfaces are not restored after a switchover or reboot.

Excluding IP Addresses from Address Pools

You can use the **ip dhcp-local excluded-address** command to specify IP addresses that you do not want the DHCP local server to supply from the default address pool. You might exclude addresses if because those addresses are already used by devices on the subnetwork.

You can exclude a single IP address or a range of addresses. To exclude a range, you specify the start-of-range IP address and the end-of-range IP address.

- To exclude a specific IP address:

```
host1(config)#ip dhcp-local excluded-address 10.10.3.4
```

- To exclude a range of IP addresses:

```
host1(config)#ip dhcp-local excluded-address 10.10.3.4 10.10.3.100
```

Configuring DHCP Local Server to Support Creation of Dynamic Subscriber Interfaces

You can use the **ip dhcp-local auto-configure agent-circuit-identifier** command to configure the DHCP local server to support the creation of dynamic subscriber interfaces built over dynamic VLANs that are based on the agent-circuit-id option (suboption 1) of the option 82 field in DHCP messages.

- Use this command within a specific virtual router context.
- This command requires that the user's DHCP control traffic and data traffic traverse the same client-facing ingress port on the E Series router.

The use of the option 82 field enables you to stack an IP interface that is associated with a particular subscriber over a dynamically created VLAN; the VLAN is dynamically created based on the agent-circuit-id option (suboption 1) that is contained in the DHCP option 82 field.

For information about configuring agent-circuit-id-based dynamic VLAN subinterfaces, see the *Configuring Dynamic Interfaces Using Bulk Configuration* chapter in the *JunosE Link Layer Configuration Guide*.

Differentiating Between Clients with the Same Client ID or Hardware Address

A JunosE Software feature enables the DHCP local server to create unique client IDs to support roaming clients and to manage situations in which two clients in the network have the same hardware address.



NOTE: This feature replaces the previous router behavior for DHCP local server client roaming and duplicate address support. The **ip dhcp-local unique-client-ids** command replaces the **ip dhcp-local inhibit-roaming** command, which has been removed from the CLI and has no effect on the DHCP local server.

You can configure the method DHCP local server uses when the router receives a DISCOVER or REQUEST packet that contains a client ID or hardware address that matches the ID or address of a currently bound client on another subnet or subinterface.

In the default configuration, the DHCP local server uses the DHCP client's subnet or subinterface to differentiate duplicate clients and support client roaming. When a new client, with a duplicate ID or hardware address, requests an address lease, DHCP assigns that client a new address and lease—the existing client's lease is unchanged.

The following table describes how the DHCP local server differentiates between a new DHCP client with the same ID or hardware address as a currently bound DHCP client.

The determination is based on whether the DHCP clients exist on the same or on different subnets and subinterfaces.

Location of DHCP Clients with Identical IDs or Addresses	How DHCP Local Server Differentiates Clients
On different subinterfaces in the same subnet	By unique subinterface
On the same subinterface in different subnets	By unique subnet
On different subinterfaces in different subnets	By unique subinterface and unique subnet
On the same subinterface in the same subnet	DHCP local server <i>cannot distinguish clients</i> with identical IDs or identical hardware addresses in this configuration

In the optional configuration, you use the **ip dhcp-local unique-client-ids** command to disable the use of the DHCP client's subnet or subinterface to differentiate between clients with duplicate client IDs or hardware addresses. When DHCP receives the request from a duplicate ID or address, DHCP terminates the address lease for the existing client and returns the address to its original address pool. DHCP then assigns a new address and lease to the new client.

We strongly recommend that you enable the **ip dhcp-local unique-client-ids** command in the following situations:

- When duplicate client IDs and duplicate hardware addresses do not exist in your network
- When the DHCP local server application interacts with DHCP relays in your network that do not support duplicate client IDs or duplicate hardware addresses

Enabling the **ip dhcp-local unique-client-ids** command in these cases enables you to properly manage DHCP clients that roam to different subnets.

The DHCP relay agent application and the DHCP relay proxy application do not support duplicate client IDs or duplicate hardware addresses.

Logging Out DHCP Local Server Subscribers

You can use the **logout subscribers** command from Privileged Exec mode to log out DHCP local server subscribers. For example, you might use this feature if you want to force a user to request a new lease or if you want to recover functional resources. The **logout subscribers** command, unlike the **clear ip address binding** command (described in [“Clearing an IP DHCP Local Server Binding” on page 490](#)), does not terminate the subscriber's user session or management representation.

This command applies to DHCP local server local-access and standalone clients, as well as to PPP users. You can log out **all** subscribers, or log out subscribers by **username**, **domain**, **virtual-router**, or **port**.

Clearing an IP DHCP Local Server Binding



NOTE: This command is deprecated and might be removed completely in a future release. The function provided by this command has been replaced by the **dhcp delete-binding** command.

You can use the **clear ip dhcp-local binding** command to force the removal of a connected user's IP address lease and associated route configuration. Using this command enables you to:

- Recover functional resources from a user who has not explicitly terminated connectivity and whose lease is unexpired.
- Discontinue connectivity to a user, prompting or forcing the user to request a new lease in order to reestablish network connectivity.

Using SNMP Traps to Monitor DHCP Local Server Events

The DHCP local server supports configurable global SNMP traps that monitor events related to the DHCP local server and local SNMP traps that are related to address pool utilization. You use the **ip dhcp-local snmpTraps** command to enable the global SNMP traps for DHCP local server.

The DHCP local server's global SNMP trap generates severity level 1 (alert), 2 (critical), and 3 (error) events. This trap helps administrators monitor DHCP local server general health, error statistics, address lease status, and protocol events. The global SNMP trap generates a severity level 4 (warning) event when a duplicate MAC address is detected. The global SNMP trap information is captured in the `dhcpLocalGeneral` logging category.

SNMP also traps events related to address pool utilization. You use the **warning** command to define the maximum and minimum threshold values and the **snmpTrap** command to generate traps when utilization occurs above or below the defined values.

For linked or shared pools, SNMP treats the members of the pool as a group, and uses the values configured for the first pool in the chain as the group's threshold.

The address pool utilization SNMP trap information is captured in the `dhcpLocalPool` logging category.



NOTE: You must configure your SNMP management client to read the MIB objects, and your SNMP trap collector must be capable of decoding the new traps. For information about setting up SNMP, see the *Configuring SNMP* chapter in the *JunosE System Basics Configuration Guide*.

Using DHCP Local Server Event Logs

To troubleshoot and monitor your DHCP local server, use the following system event logs:

- `dhcpLocalClients`—DHCP local server client events and duplicate MAC address detection
- `dhcpLocalGeneral`—DHCP local server infrastructure-related events and number of client threshold events



NOTE: The `dhcpLocalGeneral` category replaces the `dhcpLocalServerGeneral` category.

- `dhcpLocalHighAvailability`—DHCP high availability events
- `dhcpLocalPool`—DHCP local address pool events, including normal, linked, and shared pools
- `dhcpLocalProtocol`—DHCP local server protocol events

See the *JunosE System Event Logging Reference Guide* for additional information about the DHCP local server logs.

Related Documentation

- [Configuring DHCP Local Address Pools on page 491](#)
- [Configuring AAA Authentication for DHCP Local Server Standalone Mode on page 497](#)
- `clear ip dhcp-local binding`
- `dhcp delete-binding`
- `ip dhcp-local auto-configure agent-circuit-identifier`
- `ip dhcp-local excluded-address`
- `ip dhcp-local limit`
- `ip dhcp-local unique-client-ids`
- `logout subscribers`
- `service dhcp-local`
- `ipv6 local pool`

Configuring DHCP Local Address Pools

Tasks to configure DHCP local address pool include:

- [Basic Configuration of DHCP Local Address Pools on page 492](#)
- [Linking Local Address Pools on page 493](#)
- [Setting Grace Periods for Address Leases on page 494](#)

Basic Configuration of DHCP Local Address Pools

To configure the DHCP local address pool:

1. Specify the pool name and access DHCP Local Pool Configuration mode.

```
host1(config)#ip dhcp-local pool ispBoston
host1(config-dhcp-local)#
```

2. Specify the IP address of the router for the subscriber's computer to use for traffic destined for locations beyond the local subnetwork.

```
host1(config-dhcp-local)#default-router 10.10.1.1
```

The default router must be on the same subnetwork as the local server pool IP addresses that you configure with the **network** command.

You specify the IP address of a primary server, and optionally, the IP address of a secondary server.

3. (Optional) Assign a DNS server to an address pool. Some DHCP clients request the DHCP local server to assign a DNS server.

```
host1(config-dhcp-local)#dns-server 10.10.1.1
```

4. (Optional) Specify a domain name that can be returned to the subscriber if requested.

```
host1(config-dhcp-local)#domain-name ispBoston
```

The name of the domain must match the name you specified for the RADIUS vendor-specific attribute (VSA) and for authentication, authorization, accounting, and address assignment.

5. Specify the time period for which the supplied IP address is valid.

```
host1(config-dhcp-local)#lease 0 0 24
```

Specify the number of days, and optionally, the number of hours, minutes, and seconds. Use the keyword **infinite** to specify a lease that does not expire. The default lease time is 30 minutes.

6. (Optional) Link the DHCP local address pool being configured to another local address pool. See [“Linking Local Address Pools” on page 493](#) for more information about linking local address pools.

```
host1(config-dhcp-local)#link ispChicago
```

7. (Optional) Assign a NetBIOS server for subscribers. Some DHCP clients request the DHCP local server to assign a NetBIOS server.

```
host1(config-dhcp-local)#netbios-name-server 10.10.1.1 10.10.1.2
```

Specify the IP address of a primary server and, optionally, the address of a secondary server.

8. (Optional) Specify NetBIOS node type.

```
host1(config-dhcp-local)#netbios-node-type b-node
```

Specify one of the following types of NetBIOS nodes. By default, the node type is unspecified.

- **b-node**—Broadcast
 - **p-node**—Peer-to-peer
 - **m-node**—Mixed
 - **h-node**—Hybrid
9. Specify the IP addresses that the DHCP local server can provide from an address pool.

```
host1(config-dhcp-local)#network 10.10.1.0 255.255.0.0
```

Use the **force** keyword with the **no** version of the command to delete the address pool even if the pool is in use.

10. For both equal-access and standalone modes, you can reserve an IP address for a specific MAC address.

```
host1(config-dhcp-local)#reserve 10.10.13.8 0090.1a10.0552
```

11. For standalone mode, you can specify the DHCP server address that is sent to DHCP clients.

```
host1(config-dhcp-local)#server-address 10.10.20.0
```

12. (Optional) Enable Simple Network Management Protocol (SNMP) traps for local address pool utilization, including normal, linked, and shared address pools. Traps are generated based on threshold values for utilization. You can define threshold values by using the warning command. See [“Using SNMP Traps to Monitor DHCP Local Server Events” on page 490](#) for more information about SNMP and local address pools.

```
host1(config-dhcp-local)#snmpTrap
host1(config-dhcp-local)#warning 50 40
```

13. (Optional) Configure a grace period for address leases allocated from the current DHCP local address pool. Specify the number of days and, optionally, the number of hours, minutes, and seconds in the grace period.

```
host1(config-dhcp-local)#grace-period 0 12
```

This command applies only to address leases that expire. Use the **use-release-grace-period** command to also apply the configured grace period to the local pool addresses that are explicitly released by clients. See [“Setting Grace Periods for Address Leases” on page 494](#) for more information about grace periods.

14. (Optional) Specify that the grace period is applied to addresses that have been explicitly released by clients. By default, the grace period is applied only to address leases that expire, not to addresses that have been released. See [“Setting Grace Periods for Address Leases” on page 494](#) for more information about grace periods.

```
host1(config-dhcp-local)#use-release-grace-period
```

Linking Local Address Pools

In both equal-access mode and standalone mode, you can link a DHCP local pool to another local pool. The linked pool serves as a backup pool. If no addresses are available in a pool, the DHCP local server attempts to allocate an address from the linked pool. The address pools that are linked are viewed as a group.

Setting Grace Periods for Address Leases

The JunosE Software enables you to configure a grace period for a particular local address pool—the grace period is applied to all address leases associated with the address pool. The grace period is the amount of time that a client continues to retain its address lease after the lease expires or is released. An address cannot be assigned to any other client during the grace period. When the grace period expires, the address is released back to the address pool.

Grace periods help to ensure that a DHCP client retains its previously assigned IP address in situations that might normally cause a lease termination followed by a new address assignment. For example, if a client loses its lease due to a network disruption, the grace period enables the client to be reassigned the same address when the client requests an address after the network stabilizes. Grace periods are also useful during client reboots and in cases where a non-compliant or unreliable DHCP implementation triggers a lease renewal.

You configure a grace period for a local address pool. The grace period is immediately applied to all addresses that are allocated from the pool, including previously allocated addresses that are currently active—the new grace period takes precedence over a previously configured grace period for the address pool.



NOTE: Configuring a new grace period that is shorter than the address pool current grace period immediately terminates any existing address leases that are in the grace period state and that have already exceeded the length of the new grace period.

An address continues to be counted against the address pool resources while in a grace period. For example, if the address pool is exhausted, a new address cannot be assigned to other clients.

Client address leases enter the grace period in two ways—the lease might expire or the address can be explicitly released by the client. In both cases the address remains unavailable to other clients and can only be reapplied to the original client during the grace period. The address is released back to the address pool if the grace period expires before the address is reapplied to the original client.

When you configure a grace period, by default it is applied to address leases that *expire*, but not to addresses that are *released* by clients. However, you can optionally apply the grace period to released addresses.

Related Documentation

- *default-router*
- *dns-server*
- *domain-name*
- *grace-period*
- *ip dhcp-local pool*

- *lease*
- *link*
- *netbios-name-server*
- *netbios-node-type*
- *network*
- *reserve*
- *server-address*
- *snmpTrap*
- *use-release-grace-period*
- *warning*

Sending RADIUS Attributes to DHCP Subscribers Overview

The Dynamic Host Configuration Protocol (DHCP) or Dynamic Host Configuration Protocol for IPv6 (DHCPv6) server sends the following attributes returned by the RADIUS server to the DHCP or DHCPv6 subscribers on request:

- **DS-Lite-Tunnel-Name**—Offers both IPv4 and IPv6 connectivity to customers which are addressed only with an IPv6 prefix. This attribute specifies the fully qualified domain name (FQDN) of the Address Family Transition Router (AFTR) name to which the DHCPv6 client can establish an IPv4-over-IPv6 tunnel (IPv4-over-IPv6 tunnel is commonly referred to as Softwire).
- **PCP-Server-Name**—This attribute specifies the PCP server's FQDN to the DHCP Client in order to perform Port Control Protocol (PCP) functions.



NOTE:

- The DS-Lite-Tunnel-Name and PCP-Server-Name attributes are not supported when DHCP is configured in Equal-Access and Standalone mode.
- The DS-Lite-Tunnel-Name and PCP-Server-Name attributes returned by RADIUS must completely match FQDN specifications. The domain name must be expressed in a sequence of labels and the maximum length of a label is 63.

Dual-Stack Lite Tunnel Name Configuration with RADIUS and DHCPv6

The RADIUS protocol and DHCPv6 work together to return the Dual-Stack (DS)-Lite-Tunnel-Name attribute to the DHCPv6 subscribers. In this implementation, the Network Access Server (NAS) acts as a DHCPv6 server for the DHCPv6 client and a client for the RADIUS server. The following steps are performed during DHCPv6 subscriber login in authenticate mode:

- The NAS initially sends a RADIUS Access-Request message to the RADIUS server after receiving one of the following messages:
 - A DHCPv6 Solicit message from the DHCPv6 client, when an IPv6 session is used to provide connectivity to the user
 - A PPP LCP Config Request message from the DHCPv6 client, when a PPP or PPPv6 session is used to provide connectivity to the user
- The RADIUS server receives the Access-Request message, validates the NAS, and replies with an Access-Accept message if the request is approved. The Access-Accept message contains the DS-Lite-Tunnel-Name RADIUS attribute.
- The NAS uses the AFTR tunnel name returned in the DS-Lite-Tunnel-Name RADIUS attribute to populate the DHCPv6 OPTION_AFTR_NAME option in the DHCPv6 Advertise and DHCPv6 Reply messages.



NOTE: The NAS returns the AFTR tunnel name only if the DHCPv6 client has explicitly included the OPTION_AFTR_NAME option in its OPTION_ORO option.

After receiving the DS-Lite-Tunnel-Name attribute in the initial Access-Accept message, the NAS stores the received AFTR tunnel name locally. When the DHCPv6 client sends a DHCPv6 Renew message to request an extension of the lifetimes for the assigned prefix, the NAS does not initiate a new Access-Request message. The NAS retrieves the stored AFTR tunnel name and uses it in its DHCPv6 Reply message.



NOTE: The NAS sends a new Access-Request message when the DHCPv6 client initiates a Rebind/Reply message exchange with the NAS.

PCP Server Name Configuration with RADIUS and DHCP or DHCPv6

The RADIUS protocol and DHCP or DHCPv6 work together to return the PCP-Server-Name attribute to the DHCP or DHCPv6 subscribers. In this implementation, the NAS acts as a DHCP or DHCPv6 server for the DHCP or DHCPv6 client and the client for the RADIUS server. The following steps are performed during DHCP or DHCPv6 subscriber login in authenticate mode:

- The NAS initially sends a RADIUS Access-Request message to the RADIUS server after receiving one of the following messages:
 - A DHCPv6 Solicit message from the DHCPv6 client, when an IPv6 session is used to provide connectivity to the user
 - A PPP LCP Config Request message from the DHCPv6 client, when a PPP or PPPv6 session is used to provide connectivity to the user
 - A DHCP Discover message from the DHCP client, when an IP session is used to provide connectivity to the user

- The RADIUS server receives the Access-Request message, validates the NAS, and replies with an Access-Accept message if the request is approved. The Access-Accept message contains the PCP-Server-Name VSA.
- The NAS uses the PCP server name returned in the PCP-Server-Name VSA to populate the OPTION_PCP_SERVER option:
- For the DHCP client, the PCP Server Name option is populated with the PCP server name



NOTE: The NAS returns the PCP server name if the DHCP client has specified the OPTION_PCP_SERVER option in the Parameter Request List Option.

- For the DHCPv6 client, the PCP Server Name option is populated with the PCP server name



NOTE: The NAS returns the PCP server name only if the DHCPv6 client has explicitly included the OPTION_PCP_SERVER option in its OPTION_ORO option.

After receiving the PCP-Server-Name VSA in the initial Access-Accept message, the NAS stores the received PCP server name locally. When the DHCP or DHCPv6 client sends a DHCP or DHCPv6 renew message to request an extension of the lifetimes for the assigned address or prefix, the NAS does not initiate a new Access-Request message. The NAS retrieves the stored PCP server name and uses it in its reply message.



NOTE: The NAS sends a new Access-Request message when the DHCP or DHCPv6 client initiates a Rebind/Reply message exchange with the NAS.

Related Documentation

- [Juniper Networks VSAs on page 248](#)
- [RADIUS IETF Attributes on page 241](#)

Configuring AAA Authentication for DHCP Local Server Standalone Mode

The DHCP local server enables you to optionally configure AAA-based authentication of standalone mode DHCP clients. In addition to providing increased security, AAA authentication also provides RADIUS-based input to IP address pool selection for standalone mode clients. By default, clients are not authenticated in standalone mode.

Typically, an incoming DHCP client does not provide a username—therefore, the DHCP local server constructs a username based on the user's attachment parameters and optional DHCP parameters. AAA uses the constructed username to authenticate the incoming client and create the AAA subscriber record for the client. The information in the AAA subscriber record is then used to determine the IP address pool from which to

assign the address for the DHCP client. You can include the following elements in the username:

Attachment Parameters	DHCP Parameters
domain	circuit ID
user prefix	circuit type
–	MAC address
–	option 82
–	virtual router name



NOTE: The nondomain portion of a constructed username must contain at least one character. Otherwise, the DHCP local server rejects the DHCP client without performing the AAA authentication request.

When using authentication, AAA accepts the DHCP client as a subscriber—this enables you to use **show** commands to monitor configuration information and statistics about the client. You can also use the **logout subscriber** command to manage subscribers.

To configure AAA-based authentication for DHCP local server standalone mode clients:



CAUTION: Configuring authentication on the DHCP local server requires that you first disable the DHCP local server for standalone mode. Doing so removes your entire DHCP local server configuration. Therefore, if you want to configure authentication, do so before you have otherwise configured the DHCP local server.

1. Disable the DHCP local server for standalone mode.

```
host1(config)#no service dhcp-local standalone
```

2. Enable AAA-based authentication for DHCP local server standalone mode clients.

```
host1(config)#service dhcp-local standalone authenticate
```

3. Specify the password that authenticates a locally configured DHCP standalone mode client. In DHCP standalone mode, the password is presented to AAA in an authentication request.

```
host1(config)#ip dhcp-local auth password to4tooL8
```

4. Specify the domain for a username that is locally configured for a DHCP standalone mode client. The locally configured username is presented to AAA in an authentication request.

```
host1(config)#ip dhcp-local auth domain ISP1.com
```

- Specify the user-prefix for a username that is locally configured for a DHCP standalone mode client. The locally configured username is presented to AAA in an authentication request.

```
host1(config)#ip dhcp-local auth user-prefix ERX4-Boston
```

- Include optional information as part of the locally configured username for a DHCP standalone mode client. The optional information becomes part of the AAA subscriber record, and is then used to determine the IP address pool from which to assign the address for the DHCP client.

Use the following keywords to include specific information:

- circuit-identifier**—Specifies the circuit identifier of the interface on which the DHCP client's request was received.
- circuit-type**—Specifies the circuit type of the interface on which the DHCP client's request was received.
- mac-address**—Specifies the DHCP client's MAC address.
- option82**—Specifies the DHCP client's option 82 value.
- virtual-router-name**—Specifies the DHCP local server's virtual router name.

```
host1(config)#ip dhcp-local auth include virtual-router-name
host1(config)#ip dhcp-local auth include circuit-type
host1(config)#ip dhcp-local auth include circuit-identifier
```

- (Optional) Verify your authentication configuration.

```
host1(config)#show ip dhcp-local auth config
```

DHCP Local Server Authentication Configuration

```
User-Prefix           : ERX4-Boston
Domain                : ISP1.com
Password              : to4TooL8
Virtual Router        : included
Circuit Type          : included
Circuit ID            : included
MAC Address           : excluded
Option 82             : excluded
```

DHCP Local Server DHCP Options Configuration

```
RADIUS DHCP Options : excluded
```

- Related Documentation**
- ip dhcp-local auth domain*
 - ip dhcp-local auth include*
 - ip dhcp-local auth password*
 - ip dhcp-local auth user-prefix*
 - service dhcp-local*

Configuring AAA Authentication for DHCPv6 Local Server Standalone Mode

When using authentication, AAA accepts the DHCPv6 client as a subscriber—this enables you to use **show** commands to monitor configuration information and statistics about the client. You can also use the **logout subscriber** command to manage subscribers.



NOTE: The nondomain portion of a constructed username must contain at least one character. Otherwise, the DHCPv6 local server rejects the DHCPv6 client without performing the AAA authentication request.



CAUTION: Configuring authentication on the DHCPv6 local server requires that you first disable the DHCPv6 local server for standalone mode. Your entire DHCPv6 local server configuration is removed when you disable the DHCPv6 local server. Therefore, if you want to configure authentication, you must set up the authentication parameters before you configure the DHCPv6 local server for other attributes.

To configure AAA-based authentication for DHCPv6 local server standalone mode clients:

1. Disable the DHCPv6 local server for standalone mode.

```
host1(config)#no service dhcpv6-local standalone
```

2. Enable AAA-based authentication for DHCPv6 local server standalone mode clients.

```
host1(config)#service dhcpv6-local standalone authenticate
```

3. Specify the password that authenticates a locally configured DHCPv6 standalone mode client. In DHCPv6 standalone mode, the password is presented to AAA in an authentication request.

```
host1(config)#ip dhcpv6-local auth password to4tooL8
```

4. Specify the domain for a username that is locally configured for a DHCPv6 standalone mode client. The locally configured user-prefix is presented to AAA in an authentication request.

```
host1(config)#ip dhcpv6-local auth domain ISP1.com
```

5. Specify the user-prefix for a username that is locally configured for a DHCPv6 standalone mode client. The locally configured username is presented to AAA in an authentication request.

```
host1(config)#ip dhcpv6-local auth user-prefix ERX4-Boston
```

6. Include optional information as part of the locally configured username for a DHCPv6 standalone mode client. The optional information becomes part of the AAA subscriber record, and is then used to determine the IP address pool from which to assign the address for the DHCPv6 client.

Use the following keywords to include specific information:

- **circuit-identifier**—Specifies the circuit identifier of the interface on which the DHCPv6 client's request was received.
- **circuit-type**—Specifies the circuit type of the interface on which the DHCPv6 client's request was received.

```
host1(config)#ipv6 dhcpv6-local auth include circuit-identifier
host1(config)#ipv6 dhcpv6-local auth include circuit-type
```

7. (Optional) Verify your authentication configuration.

```
host1(config)#show ipv6 dhcpv6-local auth config
```

```
DHCPv6 Local Server Authentication Configuration
User-Prefix      : userPrefix
Domain           : domain
Password         : password
Circuit Type     : excluded
Circuit ID       : excluded
```

Related Documentation

- [Authentication and Accounting of IPv6 Subscribers Using the DHCPv6 Local Server Overview on page 480](#)
- [Interoperation of Authentication of IPv6 Clients and Display of Active Subscriber Information on page 482](#)
- [Monitoring DHCPv6 Local Server Authentication Information on page 596](#)
- *ipv6 dhcpv6-local auth domain*
- *ipv6 dhcpv6-local auth password*
- *ipv6 dhcpv6-local auth user-prefix*
- *service dhcpv6-local*

Configuring the DHCPv6 Local Server

In addition to the embedded DHCP local server that is used for IP version 4 (IPv4) address support, E Series routers include an embedded DHCPv6 local server. This server enables the router to function as a server for the DHCP protocol for IP version 6 (IPv6). The DHCPv6 local server sends and receives packets through IPv6 and informs IPv6 of the routing requirements of the router clients.

The DHCPv6 local server provides the following IPv6 address support:

- Delegates IPv6 prefixes to client routers; each client can have one prefix; prefixes and Domain Name System (DNS) information can be locally configured or derived from RADIUS through AAA.
- Sends a Delegated-IPv6-Prefix attribute with the PPPv6 interface name to the SRC client after delegating an IPv6 prefix to the CPE.
- Provides DNS server information to directly connected router clients.



NOTE: You must add a vendor-specific attribute to RADIUS to enable E Series routers to retrieve IPv6 DNS addresses.



NOTE: If an IPv6 prefix is not available to be delegated to requesting DHCPv6 clients, the delegating server sends the Identity Association for Prefix Delegation option, where each Identity Association for Prefix Delegation option consists of an Identity Association identifier and associated configuration information, in an Advertise message that includes a Status Code option containing the value NoPrefixAvail. For example, when a RADIUS server is used for authentication of DHCPv6 clients and the server is configured to disable the delegation of prefixes, in response to DHCPv6 Solicit messages that are received from the client, the server sends Identity Association for Prefix Delegation options in an Advertise message to the client.



NOTE: If you configure the IPv6 prefix to be delegated to the client using RADIUS attributes and the client is authenticated by the RADIUS server, it is not necessary to configure the prefix associated with it using this procedure. In such a scenario, the prefix derives its settings from the Framed-IPv6-Prefix [23] or the Delegated-IPv6-Prefix [97] RADIUS IETF attributes.



NOTE: The DHCPv6 local server supports the DHCPv6 Rapid Commit option (DHCPv6 option 14). When a DHCPv6 client includes the DHCPv6 Rapid Commit option in the Solicit messages, the server recognizes this option and the client terminates the waiting process as soon as it receives a Reply message with the Rapid Commit option from the server. The server and client then use a two-message exchange (Solicit and Reply) to configure clients, rather than the default four-method exchange (Solicit, Advertise, Request, and Reply). The two-message exchange provides faster client configuration, and is beneficial in environments in which networks are under a heavy load.

Use the following steps to configure the DHCPv6 local server:

1. Enable the DHCPv6 local server.

```
host1(config)#service dhcpv6-local
```

2. Specify the IPv6 prefix and lifetime that are to be delegated to the DHCPv6 client. The specified prefix is delegated by the DHCPv6 local server when requested by the client.

```
host1(config-if)#ipv6 dhcpv6-local delegated-prefix 2001:db8:17::/48 lifetime infinite
```

Use the **lifetime** keyword to specify the time period for which the prefix is valid. This lifetime overrides the default lifetime that is set in Global Configuration mode. If no lifetime is specified, the default lifetime is assigned.

- Specify the number of days and, optionally, the number of hours, minutes, and seconds. You cannot specify a lifetime of zero (that is, you cannot set the days, hours, minutes, and seconds fields all to zero).
 - Use the keyword **infinite** to specify a lifetime that does not expire.
3. Specify the name of a DNS domain for DHCPv6 clients in the current virtual router to search. You can specify a maximum of four DNS domains for a DHCPv6 local server's search list.

```
host1(config)#ipv6 dhcpv6-local dns-domain-search xyzcorporation.com
host1(config)#ipv6 dhcpv6-local dns-domain-search xyzcorp.com
```

4. Specify the IPv6 address of the DNS server and to assign the server to the DHCPv6 clients in the current virtual router. You can specify a maximum of four DNS servers.

```
host1(config)#ipv6 dhcpv6-local dns-server 2001:db8:18::
```

5. Set the default lifetime for which a prefix delegated by this DHCPv6 local server is valid. This default is overridden by an interface-specific lifetime.

```
host1(config)#ipv6 dhcpv6-local prefix-lifetime infinite
```

- Specify the number of days and, optionally, the number of hours, minutes, and seconds. You cannot specify a lifetime of zero (that is, you cannot set the days, hours, minutes, and seconds fields all to zero).
 - Use the keyword **infinite** to specify a lifetime that does not expire.
6. Specify the DHCP unique identifier (DUID) type to be used in the communication between the DHCPv6 local server and clients. You can configure the type of DUID to be either Type 2 or Type 3. These two types are currently supported by the DHCPv6 local server application in JunosE Software. The Type 1 DUID is not supported by JunosE Software.

```
host1(config)#ipv6 dhcpv6-local duid-type 3
```

7. (Optional) Specify the maximum number of IPv6 prefixes that the DHCPv6 local server can delegate to each ATM interface, Ethernet subnet, VLAN interface, or POS access interface, or to a particular interface of subinterface. For more information about the limiting of delegated IPv6 prefixes, see ["Limiting the Maximum Number of IPv6 Prefixes Delegated Per Interface by the DHCPv6 Local Server Overview"](#) on page 483.

- To set the maximum number of IPv6 prefixes globally for an interface type:

```
host1(config)#ipv6 dhcpv6-local limit vlan 300
```

- To set the maximum number of IPv6 prefixes for a particular interface:

```
host1(config)#ipv6 dhcpv6-local limit interface gigabitEthernet 6/0 500
```

8. (Optional) Enable the DHCPv6 local server to send a Delegated-IPv6-Prefix with the PPPv6 interface name to the SRC client after delegating the IPv6 prefix to the CPE. The SRC client sends the Delegated-IPv6-Prefix to an SRC server.

By default, the DHCPv6 local server does not send the Delegated-IPv6-Prefix to the SRC client. For more information about this feature, see [“Sending of Delegated-IPv6-Prefix from the DHCPv6 Local Server to SRC Client Overview” on page 56.](#)

```
host1(config)#sscc enable cops-pr
host1(config)# sscc protocol ipv6
host1(config)#ipv6 dhcpv6-local sscc enable
```

- | | |
|------------------------------|--|
| Related Documentation | <ul style="list-style-type: none">• Monitoring the DHCPv6 Local Server to Send Delegated-IPv6-Prefix to SRC Client on page 602• Monitoring the Maximum Number of IPv6 Prefixes Delegated Per Interface by the DHCPv6 Local Server on page 600• <i>ip dhcp-local auth domain</i>• <i>ipv6 dhcpv6-local delegated-prefix</i>• <i>ipv6 dhcpv6-local dns-domain-search</i>• <i>ipv6 dhcpv6-local dns-server</i>• <i>ipv6 dhcpv6-local duid-type</i>• <i>ipv6 dhcpv6-local limit</i>• <i>ipv6 dhcpv6-local prefix-lifetime</i>• <i>ipv6 dhcpv6-local sscc</i>• <i>service dhcpv6-local</i>• <i>sscc enable</i>• <i>sscc protocol ipv6</i> |
|------------------------------|--|

Configuring the Type of DHCP Unique ID for DHCPv6 Local Servers

You can configure the type of DHCP unique identifier (DUID) using the **ipv6 dhcpv6-local duid-type** *duidType* command in Global Configuration mode to be either Type 2 or Type 3. These two types are currently supported by the DHCPv6 local server application in JunosE Software. The Type 1 DUID is not supported by the DHCPv6 local server in JunosE Software. However, DHCPv6 clients support DUID Types 1, 2, and 3.

To configure the DUID type:

1. Enable the DHCPv6 local server.

```
host1(config)#service dhcpv6-local
```
2. Specify the DUID type to be used during the identity verification of the server and the client.


```
host1(config)#ipv6 dhcpv6-local duid 3
```

In this example, the DUID type is set as Type 3, which is used by devices that have a permanently connected network interface with a link-layer address, and do not have a nonvolatile, writable stable storage.



NOTE: You must enable the DHCPv6 local server using the **service dhcpv6-local** command before configuring the DUID type. Otherwise, an error message states that the DHCPv6 local server is not configured on the router.

Related Documentation

- [DHCP Unique ID for Clients and Servers Overview on page 478](#)
- `ipv6 dhcpv6-local duid-type`

Deleting DHCPv6 Client Bindings

The JunosE Software enables you to manage your router's DHCPv6 local server client bindings. The client binding associates an IPv6 prefix with a unique DHCP ID (DUID) of the subscriber client. To view information about current DHCPv6 client bindings and track lease times of a specific client binding, use the **show ipv6 dhcpv6-local binding** command.

To delete a client binding and the associated route configuration when the DHCPv6 client binding is no longer needed, use the **dhcpv6 delete-binding** command. You can delete the DHCPv6 client bindings instead of waiting for the lease timer to expire. Use the following keywords and variables with the **dhcpv6 delete-binding** command to specify (filter) the client bindings you want to delete:

- **all**—All DHCPv6 local server client bindings
- *ipv6Prefix*—IPv6 prefix (address and subnetwork mask) of the DHCPv6 clients; for example, 2002:2:4:1::/64
- *string*—Local address pool name; for example, server4pool



NOTE: After a stateful SRP switchover, in a scaled environment, the interface strings associated with DHCPv6 client bindings might not be displayed in the output of the **show** commands used to view information about client bindings if you issue the **show** command immediately after a stateful SRP switchover. These **show** commands display interface strings in the output only if the restoration of IPv6 interfaces on the router is complete after the SRP warm switchover. After the restoration of IPv6 interfaces is complete, interface strings are displayed properly in the output of the **show** commands available for this purpose.

You can remove all DHCPv6 client bindings, all DHCPv6 client bindings of a particular type, or a specified DHCPv6 client binding that meets the deletion criteria you specify.

- To delete all DHCPv6 client bindings on virtual router vr1:

```
host1:vr1#dhcpv6 delete-binding all
```

- To delete DHCPv6 client bindings with the specified IPv6 prefix:

```
host1:vr1#dhcpv6 delete-binding 2002:2:4:1::/64
```

- To delete a group of DHCPv6 client bindings that were assigned prefix from the local pool:

```
host1:vr2#dhcpv6 delete-binding server4pool
```

The router does not notify the DHCPv6 client when you use the **dhcpv6 delete-binding** command. To verify that the DHCPv6 client bindings have been deleted, use the **show ipv6 dhcpv6-local binding** command.

In JunosE Release 11.3.0, when DHCPv6 client bindings are brought up over a PPPv6 session, on a router that acts as an L2TP network server (LNS) and is enabled for stateful line module switchover, the client bindings are removed when the primary line module fails and the spare line module takes over as the primary. This behavior occurs because the underlying dynamic IPv6 over PPP interface goes down temporarily (when the subscriber session is disrupted briefly) before the interface becomes operational again on the newly active primary module. When the dynamic IPv6 over PPP interface goes down temporarily (when the stateful switchover process is in progress), the DHCPv6 client binding and the access route for that interface are deleted. Similarly, DHCPv6 bindings are deleted when a PPP subscriber logs out and then back in. In such scenarios, the client needs to send a renew or rebind request to the DHCP server to enable the DHCPv6 binding to be re-created.

Beginning with JunosE Release 12.0.0, DHCPv6 client bindings and access routes that are created over a PPPv6 session on an LNS device enabled for stateful line module switchover are retained when the dynamic IPv6 over PPP interface temporarily goes down during the stateful switchover operation. DHCPv6 client bindings and the associated route configuration are deleted only when the interface is deleted and not during the interface down event.

DHCPv6 client bindings and access routes that are created over a PPPv6 session, on a router that acts as an LNS and is enabled for stateful line module switchover, are retained when the dynamic IPv6 over PPP interface goes down temporarily during the stateful switchover operation. When the stateful switchover procedure is complete, the interface is re-created on the newly active primary module and the DHCPv6 bindings are also retained. The same behavior of preservation of DHCPv6 bindings is applicable when a PPPv6 subscriber logs out and then back in.

Related Documentation

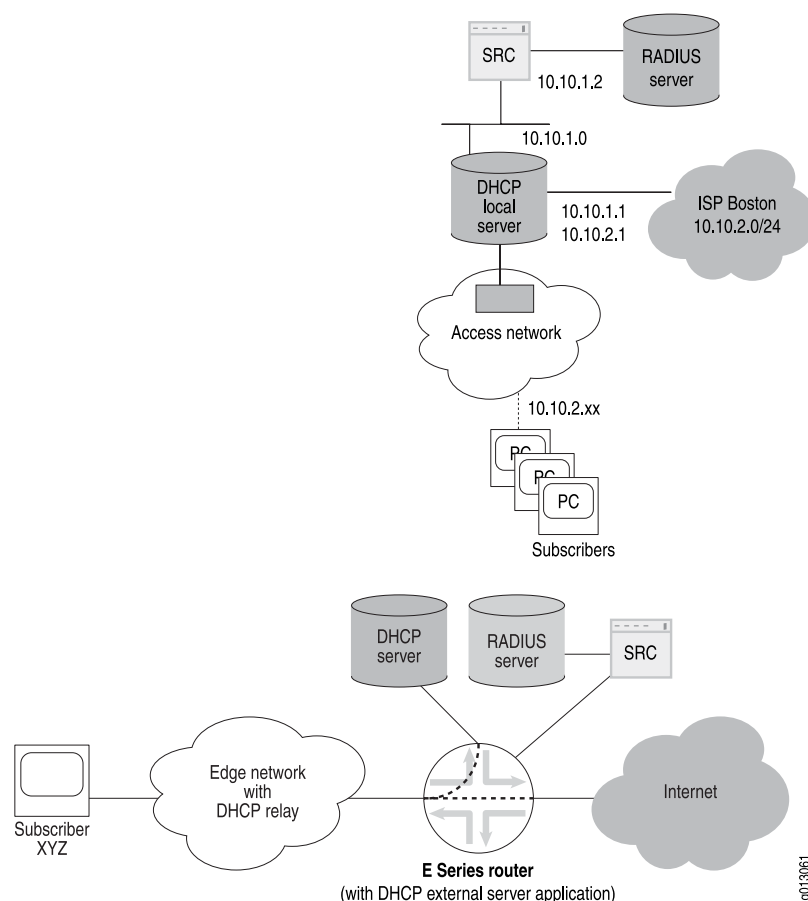
- [Monitoring DHCPv6 Local Server Binding Information on page 593](#)
- *dhcpv6 delete-binding*
- *show ipv6 dhcpv6-local binding*

Configuring the Router to Work with the SRC Software

E Series Broadband Services Routers have an embedded SRC client that interacts with the SRC software. For information about configuring the SRC client, see [“SRC Client Configuration Overview” on page 47](#).

Figure 12 on page 507 shows the scenario for this example. Subscribers obtain access to ISP Boston via a router. Subscribers log in through the SRC software, and a RADIUS server provides authentication.

Figure 12: Non-PPP Equal-Access Configuration Example



The following steps describe how to configure this scenario.

1. Configure interfaces on the router.

```
host1(config)#interface loopback 0
host1(config-if)#ip address 10.10.1.1 255.255.255.255
host1(config-if)#ip address 10.10.2.1 255.255.255.255 secondary
host1(config-if)#exit
host1(config)#interface fastEthernet 2/0
host1(config-if)#ip unnumbered loopback 0
```

2. Configure the parameters to enable the router to forward authentication requests to the RADIUS server.

```
host1(config)#radius authentication server 10.10.1.2
host1(config)#udp-port 1645
host1(config)#key radius
```

3. Specify the authentication method.

```
host1(config)#aaa authentication ppp default radius
```

Or

```
host1(config)#aaa authentication ppp default none
```

4. Enable the DHCP local server.

```
host1(config)#service dhcp-local
```

5. Specify the IP addresses that are in use, so that the DHCP local server cannot assign these addresses.

```
host1(config)#ip dhcp-local excluded-address 10.10.1.1
host1(config)#ip dhcp-local excluded-address 10.10.1.2
```

6. Configure the DHCP local server to provide IP addresses to subscribers of ISP Boston.

```
host1(config)#ip dhcp-local pool ispBoston
host1(config-dhcp-local)#network 10.10.2.0 255.255.255.0
host1(config-dhcp-local)#domain-name ispBoston
host1(config-dhcp-local)#default-router 10.10.2.1
host1(config-dhcp-local)#lease 0 0 10
host1(config-dhcp-local)#ip dhcp-local limit atm 5
```

7. Configure the SRC client.

```
host1(config)#sscc enable
host1(config)#sscc retryTimer 200
host1(config)#sscc primary address 10.10.1.2 port 3288
```

Related Documentation

- [Configuring the SRC Client on page 97](#)
- *aaa authentication default*
- *default-router*
- *domain-name*
- *interface fastEthernet*
- *interface loopback*
- *ip address*
- *ip dhcp-local excluded-address*
- *ip dhcp-local limit*
- *ip dhcp-local pool*
- *ip unnumbered*
- *key*

- *lease*
- *network*
- *radius authentication server*
- *sscc address*
- *sscc enable*
- *sscc retryTimer*
- *udp-port*

CHAPTER 21

Configuring DHCP Relay

The Dynamic Host Configuration Protocol (DHCP) provides a mechanism through which computers using Transmission Control Protocol/IP (TCP/IP) can obtain protocol configuration parameters automatically from a DHCP server on the network.

The following sections describe how to configure your E Series router to provide DHCP support:

- [DHCP Relay and BOOTP Relay Overview on page 512](#)
- [Configuring DHCP Relay Settings on page 512](#)
- [Using the Broadcast Flag Setting to Control Transmission of DHCP Reply Packets on page 516](#)
- [Interaction of DHCP Relay Broadcast Flag with the Layer 2 Unicast Transmission Method on page 516](#)
- [Preventing DHCP Relay from Installing Host Routes by Default on page 518](#)
- [Example: Preventing Installation of Host Routes on page 519](#)
- [Including Relay Agent Option Values in the PPPoE Remote Circuit ID on page 520](#)
- [Configuring the Layer 2 Unicast Transmission Method for Reply Packets to DHCP Clients on page 521](#)
- [Using Option 60 Strings to Forward Client Traffic to Specific DHCP Servers on page 522](#)
- [Configuring the DHCP Relay Option 60 Attribute for Traffic Forwarding on page 524](#)
- [Relaying DHCP Packets That Originate from a Cable Modem on page 525](#)
- [DHCP Relay Agent Information Option \(Option 82\) Suboption Values Overview on page 525](#)
- [Configuring the DHCP Relay Agent Option 82 Settings on page 529](#)
- [Using the set dhcp relay agent sub-option Command to Enable Option 82 Suboption Support on page 531](#)
- [Using the set dhcp relay agent Command to Enable Option 82 Suboption Support on page 533](#)
- [Example: Using DHCP Relay Option 82 to Pass IEEE 802.1p Values to DHCP Servers on page 535](#)
- [Rate of DHCP Client Packets Processed by DHCP Relay Overview on page 539](#)
- [Configuring the Rate of Client Packets Processed by DHCP Relay on page 540](#)

- [Configuring DHCP Relay Proxy on page 540](#)
- [Managing Host Routes Using DHCP Relay Proxy on page 542](#)

DHCP Relay and BOOTP Relay Overview

The DHCP relay feature relays a request from a remote client to a DHCP server for an IP address. When the router receives a DHCP request from an IP client, it forwards the request to the DHCP server and passes the response back to the IP client.

Configuring DHCP relay also enables bootstrap protocol (BOOTP) relay. The router relays any BOOTP requests it receives to the same set of servers that you configured for DHCP relay. A DHCP server can respond to the BOOTP request only if it is also a BOOTP server. The router relays any BOOTP responses it receives to the originator of the BOOTP request. If you do not configure DHCP relay, then BOOTP relay is disabled.

The router must wait for an acknowledgment from the DHCP server that the assigned address has been accepted. The IP client must accept an IP address from one of the servers. When the DHCP server sends an acknowledgment message back to the DHCP client via the router, the router updates its routing table with the IP address of the client.

If a DHCP relay request is received on an unnumbered interface, the router determines the loopback address for that interface and passes that IP address to the server.

DHCP carries other important configuration parameters, such as the subnet mask, default router, and DNS server. You can also use the DHCP relay agent information option (option 82) to add information to the DHCP packets sent to DHCP servers—the additional information, in the form of suboptions to the option 82 value, helps you to manage the IP address and service level assignments granted to your subscribers. For example, you can add the E Series hostname or the virtual router name to the front of the Agent Circuit ID suboption (suboption 1) of the DHCP relay agent information option (option 82). See [“Configuring Relay Agent Information Option \(Option 82\) Suboption Values” on page 530](#).

Related Documentation

- [Configuring DHCP Relay Settings on page 512](#)
- [Enabling DHCP Relay on page 513](#)
- [Configuring DHCP Relay Proxy on page 540](#)
- [Using the Broadcast Flag Setting to Control Transmission of DHCP Reply Packets on page 516](#)

Configuring DHCP Relay Settings

You can configure DHCP relay using the following set of tasks:

- [Enabling DHCP Relay on page 513](#)
- [Removing Access Routes from Routing Tables and NVS on page 513](#)
- [Treating All Packets as Originating at Trusted Sources on page 514](#)
- [Assigning the Giaddr to Source IP Address on page 514](#)

- [Protecting Against Spoofed Giaddr and Relay Agent Option Values on page 514](#)
- [Using the Giaddr to Identify the Primary Interface for Dynamic Subscriber Interfaces on page 515](#)

Enabling DHCP Relay

You use the **set dhcp relay** command to create and enable DHCP relay in the current virtual router.

- Include the IP address variable to enable DHCP relay and BOOTP relay and to specify an IP address for the DHCP server. When you include the IP address of a DHCP server, the router adds the IP address to the list of DHCP servers (up to five) and forwards all request packets to all configured servers.

Issuing this command also enables relay of BOOTP requests to the configured DHCP servers. If one of the DHCP servers is also a BOOTP server and responds, the router relays the response to the request originator.

```
host1(config)#set dhcp relay 192.168.29.10
```

- Use this command without an IP address to create the DHCP relay independent of any DHCP servers. Use this version of the command when configuring support for DHCP vendor-option strings (option 60). For information about configuring option 60 support, see [“Using Option 60 Strings to Forward Client Traffic to Specific DHCP Servers” on page 522](#).

```
host1(config)#set dhcp relay
```

- Use the **no** version with an IP address to remove the specified DHCP server:

```
host1(config)#no set dhcp relay 192.168.29.25
```

Removing Access Routes from Routing Tables and NVS

You can remove existing access routes for an interface from routing tables and nonvolatile storage (NVS).

This command removes all installed host routes from IP and deletes host routes from mirrored storage and NVS for specified interfaces. In relay proxy mode, this command enforces consistent state of the route and client database and discards all client information for specified interfaces.

Because DHCP relay cannot distinguish between temporary dynamic interface deletions—where the interface is subsequently re-created—and permanent deletions, sometimes it retains routing information for dynamic interfaces that have already been deleted. You can use the **unknown** keyword with the **dhcp relay discard access-routes** command to remove the routing information for these interfaces.

- To remove access routes:

```
host1(config)#set dhcp relay discard-access-routes
```



NOTE: When this feature is configured, the client bypasses the DHCP relay component and communicates directly with the DHCP server to request address renewal or to release the address. The DHCP relay component has no role in determining when or whether to remove the installed host route.

Treating All Packets as Originating at Trusted Sources

By default, the DHCP relay treats all packets destined for DHCP servers as if the packets originated at an untrusted source; if the packets have a gateway IP address (giaddr) of 0 and if option 82 information is present, these packets are dropped.

- To enable the trust-all method on the DHCP relay:

```
host1(config)#set dhcp relay trust-all
```

In the trust-all method, the DHCP relay treats the packets as if they are from trusted sources and forwards the packets to the DHCP server. When you enable this command:

- If the DHCP packets contain option 82 and a giaddr field of 0, the DHCP relay inserts its giaddr into the packets and then forwards the packets.
- If the DHCP relay is configured to add option 82, it does not add an additional option 82 if one is already present in the DHCP packets.

Assigning the Giaddr to Source IP Address

As a security measure, DHCP servers typically use the giaddr included in DHCP packets to ensure that the packets come from a recognized DHCP gateway. The servers verify that the giaddr in the DHCP packet matches the source IP address in the IP packet header. You can use the **set dhcp relay assign-giaddr-source-ip** command to specify that the DHCP relay and DHCP relay proxy assign the giaddr to the source IP packet header of packets they send to DHCP servers—the DHCP servers can then compare the giaddr in the IP packet header to the giaddr in the DHCP packets.

- To assign the giaddr to the source IP packet header:

```
host1(config)#set dhcp relay assign-giaddr-source-ip
```

Protecting Against Spoofed Giaddr and Relay Agent Option Values

DHCP relay includes an override feature that provides enhanced security to protect against spoofed giaddr and relay agent option (option 82) values in packets destined for DHCP servers.

DHCP relay can detect spoofed giaddrs when the giaddr value is equal to a local IP address on which the DHCP relay can be accessed; otherwise, DHCP relay does not detect spoofed giaddrs. Also, DHCP relay does not detect spoofed relay agent option values.

Spoofed giaddrs are a concern when the DHCP relay is used if the giaddr value in received DHCP packets is different from the local IP address on which the DHCP relay is accessed. In this situation, DHCP relay always honors the giaddr. To configure DHCP relay to override

all giaddrs (including valid giaddrs) that are received from downstream network elements, use the **set dhcp relay override** command with the **giaddr** keyword. DHCP relay then takes control of the client, adding its own giaddr to the packets before forwarding the packets to the DHCP server.

Spoofed relay agent options are a concern if the giaddr is not null, or if it is null and the DHCP relay is operating in the trust-all method. In these two situations, DHCP relay always honors the relay agent option value in received DHCP packets.

- To protect against spoofed giaddrs and relay agent option values:

```
host1(config)#set dhcp relay override agent-option
```

DHCP relay then overrides all relay agent option values that are received from downstream network elements, performing one of the following actions:

- If the DHCP relay is configured to add relay agent option 82 to the packets, it clears the existing option 82 values and inserts the new values.
- If the DHCP relay is not configured to add relay agent option 82, it clears the existing option values but does not add any new values.

Using the Giaddr to Identify the Primary Interface for Dynamic Subscriber Interfaces

When creating dynamic subscriber interfaces, the router builds the dynamic interfaces on the associated primary interface. By default, the router identifies the primary interface based on the interface on which DHCP client discover packets are received. The router then builds all dynamic interfaces on that primary interface.

In some cases you might want more control over the determination of the primary interface and you might not want to use the primary interface that is determined by the default behavior. The JunosE Software enables you to configure DHCP relay to use information in the giaddr in DHCP ACK messages to specify which interface is to be used as the primary interface. This capability allows you to build dynamic interfaces on the primary interface of your choice.

- To use information in the giaddr to identify the primary interface for dynamic subscriber interfaces:

```
host1(config)#set dhcp relay giaddr-selects-interface
```

Related Documentation

- [Configuring DHCP Relay Proxy on page 540](#)
- [Using the Broadcast Flag Setting to Control Transmission of DHCP Reply Packets on page 516](#)
- *set dhcp relay*
- *set dhcp relay assign-giaddr-source-ip*
- *set dhcp relay giaddr-selects-interface*
- *set dhcp relay override*
- *set dhcp relay trust-all*

Using the Broadcast Flag Setting to Control Transmission of DHCP Reply Packets

Each DHCP request packet includes a broadcast flag that, if set, specifies how to transmit DHCP Offer reply packets and DHCP ACK and NAK reply packets to DHCP clients during the discovery process. To configure DHCP relay and DHCP relay proxy to use the setting of the broadcast flag to control the transmission of DHCP Offer, DHCP ACK, and DHCP NAK reply packets, use the **set dhcp relay broadcast-flag-replies** command from Global Configuration mode.

host1(config)#set dhcp relay broadcast-flag-replies

When you issue the **set dhcp relay broadcast-flag-replies** command, the method that DHCP relay and DHCP relay proxy use to transmit DHCP Offer reply packets and ACK and NAK reply packets depends on whether the broadcast flag in the DHCP request packet is set or not set, as follows:

- If the broadcast flag is set in the DHCP request packet, using the **set dhcp relay broadcast-flag-replies** command causes DHCP relay and DHCP relay proxy to broadcast DHCP reply packets to clients.
- If the broadcast flag is not set in the DHCP request packet, using the **set dhcp relay broadcast-flag-replies** command causes DHCP relay and DHCP relay proxy to use the layer 2 unicast transmission method to send DHCP reply packets using the client's layer 2 (MAC) address and layer 3 (IP) unicast address.

There are exceptions to this behavior for DHCP relay proxy when the DHCP client is already bound to an IP address or is renewing the lease on its IP address. For information, see the *Behavior for Bound Clients and Address Renewals* section in [“Managing Host Routes Using DHCP Relay Proxy” on page 542](#).

To display whether support for broadcast flag replies is currently on or off on the router, use the **show dhcp relay** command. For information, see *Monitoring and Troubleshooting DHCP*.

To troubleshoot applications that use this feature, you can use the dhcpCapture system event log category. For information about how to log system events, see *JunosE System Event Logging Reference Guide*.

Related Documentation

- [Interaction of DHCP Relay Broadcast Flag with the Layer 2 Unicast Transmission Method on page 516](#)
- [Configuring the Layer 2 Unicast Transmission Method for Reply Packets to DHCP Clients on page 521](#)
- `set dhcp relay broadcast-flag-replies`

Interaction of DHCP Relay Broadcast Flag with the Layer 2 Unicast Transmission Method

As described in [“Configuring the Layer 2 Unicast Transmission Method for Reply Packets to DHCP Clients” on page 521](#), you can use the **set dhcp relay layer2-unicast-replies**

command to configure DHCP relay and DHCP relay proxy to use the layer 2 unicast and layer 3 broadcast transmission method to send DHCP Offer reply packets and DHCP ACK and NAK reply packets to clients.

The **set dhcp relay broadcast-flag-replies** command and the **set dhcp relay layer2-unicast-replies** command are mutually exclusive. If you attempt to issue the **set dhcp relay broadcast-flag-replies** command when the **set dhcp relay layer2-unicast-replies** command is already in effect, the operation fails and the router displays the following message:

```
% layer2-unicast-replies and broadcast-flag-replies are mutually exclusive
```

If this message appears, you must first issue the **no set dhcp relay layer2-unicast-replies** command to disable layer 2 unicast replies, and then issue the **set dhcp relay broadcast-flag-replies** command again to enable broadcast flag replies.

Table 110 on page 517 summarizes how the configuration of the **set dhcp relay broadcast-flag-replies** command and the **set dhcp relay layer2-unicast-replies** command interacts with the setting of the broadcast flag in DHCP request packets to control how the router transmits DHCP reply packets to clients during the discovery process. Because these commands are mutually exclusive, broadcast flag replies and layer 2 unicast replies cannot both be enabled on the router at the same time.

Table 110: Router Configuration and Transmission of DHCP Reply Packets

Broadcast Flag Replies	Layer 2 Unicast Replies	Router Behavior if Broadcast Flag Set	Router Behavior if Broadcast Flag Not Set
Enabled (on)	Disabled (off)	DHCP relay and DHCP relay proxy broadcast DHCP reply packets to clients.	DHCP relay and DHCP relay proxy use layer 2 unicast and layer 3 unicast transmission to send DHCP reply packets to clients.
Disabled (off)	Enabled (on)	DHCP relay and DHCP relay proxy use layer 2 unicast and layer 3 broadcast transmission to send DHCP reply packets to clients.	DHCP relay and DHCP relay proxy use layer 2 unicast and layer 3 broadcast transmission to send DHCP reply packets to clients.
Disabled (off)	Disabled (off)	DHCP relay and DHCP relay proxy broadcast DHCP reply packets to clients. For information about exceptions to this behavior for DHCP relay proxy, see the <i>Behavior for Bound Clients and Address Renewals</i> section in “Managing Host Routes Using DHCP Relay Proxy” on page 542.	DHCP relay and DHCP relay proxy broadcast DHCP reply packets to clients. For information about exceptions to this behavior for DHCP relay proxy, see the <i>Behavior for Bound Clients and Address Renewals</i> section in “Managing Host Routes Using DHCP Relay Proxy” on page 542.

- Related Documentation**
- [Using the Broadcast Flag Setting to Control Transmission of DHCP Reply Packets on page 516](#)
 - [Configuring the Layer 2 Unicast Transmission Method for Reply Packets to DHCP Clients on page 521](#)
 - *set dhcp relay broadcast-flag-replies*
 - *set dhcp relay layer2-unicast-replies*

Preventing DHCP Relay from Installing Host Routes by Default

The Address Resolution Protocol (ARP) performs spoof checking on all incoming ARP requests by default. For each incoming packet, ARP does a route lookup on the source IP address to determine the interface on which that IP address was routed. ARP then verifies that the interface on which the packet was received matches the routed interface. If the interface on which the packet was received does not match the routed interface, the router drops the packet.

When you configure applications such as DHCP relay that automatically install routes, you must ensure that the routes are correctly installed for your configuration. DHCP relay installs host routes by default, which is required in certain configurations to enable address renewals from the DHCP server to work properly. However, the default installation of host routes might cause a conflict when you configure DHCP relay with static subscriber interfaces. To avoid these configuration conflicts, use the **set dhcp relay inhibit-access-route-creation** command to prevent DHCP relay from installing host routes by default. The command enforces consistent state of the route and client database.

In relay mode, this command removes all installed host routes from IP, deletes all host routes from mirrored storage and NVS, and stops accumulating host route information.

In relay proxy mode, this command removes all installed host routes from IP, deletes all NVS client data, and stops installing host routes for newly bound clients in IP. However, it does preserve the client data in mirrored storage and continues preservation of newly bound clients in mirrored storage.

The **no set dhcp relay inhibit-access-route-creation** command enforces consistent state of the route and client database. In relay proxy mode, after the unified ISSU is completed and normal operations resume, this command installs a host route for all existing bound clients in IP and saves it in NVS.

- Related Documentation**
- [Example: Preventing Installation of Host Routes on page 519](#)
 - *set dhcp relay*

Example: Preventing Installation of Host Routes

This example describes the configuration of a DHCP relay to prevent installation of host routes.

- [Requirements on page 519](#)
- [Overview on page 519](#)
- [Configuring DHCP Relay to Prevent Installation of Host Routes on page 519](#)

Requirements

This example uses the following software and hardware components:

- JunosE Release 7.1.0 or higher-numbered releases
- E Series router (ERX7xx models, ERX14xx models, the ERX310 router, the E120 router, or the E320 router)
- ASIC-based line modules that support Fast Ethernet or Gigabit Ethernet

Overview

This example describes a sample procedure for configuring multiple subscribers over a particular static subscriber interface (ip53001 in this example)—you might use commands similar to the following to create demultiplexer table entries and a subnet route that points to the static subscriber interface.

In the example, the host routes are associated with the primary IP interface on Gigabit Ethernet 1/0. Because the host routes are statically configured with the subscriber interface, there is no need for the router to install DHCP host routes. Therefore, in step 7, the `set dhcp relay inhibit-access-route-creation` command is used to prevent DHCP relay from installing host routes.

In the example, if you do not prevent DHCP relay from installing host routes, the ARP spoof-checking mechanism associates the ARP traffic with the primary IP interface (Gigabit Ethernet 1/0), although packets actually arrive on the subscriber interface (ip53001), causing the router to detect a spoof and drop the packet.

Configuring DHCP Relay to Prevent Installation of Host Routes

Step-by-Step Procedure

This example describes how you can configure the DHCP relay to not install host routes:

1. Create a shared IP interface.
`host1(config)#interface ip ip53001`
2. Associate the shared IP interface with a static layer 2 interface.
`host1(config-if)#ip share-interface gigabitEthernet 1/0`
3. Make the shared interface an unnumbered interface.
`host1(config-if)#ip unnumbered loopback 53`

4. Specify the source addresses that the subscriber interface uses to demultiplex traffic.

```
host1(config-if)#ip source-prefix 10.10.10.0 255.255.255.252
```

5. Exit Interface Configuration mode.

```
host1(config-if)#exit
```

6. Create a static route that sends traffic for destination address 10.10.10.0 to subscriber interface ip53001.

```
host1(config)#ip route 10.10.10.0 255.255.255.252 ip ip53001
```

7. Prevent DHCP relay from installing host routes—this avoids a conflict that can cause undesirable ARP behavior.

```
host1(config)#set dhcp relay inhibit-access-route-creation
```

Related Documentation

- [Preventing DHCP Relay from Installing Host Routes by Default on page 518](#)
- *interface ip*
- *ip route*
- *ip share-interface*
- *ip source-prefix*
- *ip unnumbered*
- *set dhcp relay*

Including Relay Agent Option Values in the PPPoE Remote Circuit ID

You can enable the router to capture and format a vendor-specific tag containing a PPPoE remote circuit ID value transmitted from a digital subscriber line access multiplexer (DSLAM) device. The router can then send this value to a Remote Authentication Dial-In User Service (RADIUS) server or to a Layer 2 Tunneling Protocol (L2TP) network server (LNS) to uniquely identify subscriber locations.

By default, the router formats the captured PPPoE remote circuit ID to include only the agent-circuit-id suboption (suboption 1) of the DHCP relay agent information option (option 82). You can use the **radius remote-circuit-id-format** command to configure the following nondefault formats for the PPPoE remote circuit ID value:

- Include either or both of the agent-circuit-id (suboption 1) and agent-remote-id (suboption 2) suboptions of the DHCP relay agent information option, with or without the NAS-Identifier [32] RADIUS attribute.
- Append the agent-circuit-id suboption value to an interface specifier that is consistent with the recommended format in the DSL Forum Technical Report (TR)-101—Migration to Ethernet-Based DSL Aggregation (April 2006).

For information about configuring the PPPoE remote circuit ID, see the *Using the PPPoE Remote Circuit ID to Identify Subscribers* and *Configuring PPPoE Remote Circuit ID Capture* sections in *JunosE Link Layer Configuration Guide*.

Related Documentation

- *radius remote-circuit-id-format*

Configuring the Layer 2 Unicast Transmission Method for Reply Packets to DHCP Clients

By default, DHCP relay and relay proxy broadcast DHCP Offer reply packets and DHCP ACK and NAK reply packets to DHCP clients during the discovery process. In some environments, this default broadcast method might be a security concern because all clients can receive packets intended for all other clients.

You use the **set dhcp relay layer2-unicast-replies** command in Global Configuration mode to configure the optional layer 2 unicast and layer 3 broadcast transmission method for DHCP relay and DHCP relay proxy. This method uses the client's layer 2 (MAC) address and layer 3 (IP) broadcast address to provide secure transmission of DHCP Offer reply packets and ACK and NAK reply packets. The optional layer 2 unicast method enables reply packets to be broadcast through the layer 3 network but received only by the specified client.

There are exceptions to this behavior for DHCP relay proxy when the DHCP client is already bound to an IP address or is renewing the lease on its IP address. For information, see the *Behavior for Bound Clients and Address Renewals* section in [“Managing Host Routes Using DHCP Relay Proxy” on page 542](#).

To display whether the layer 2 unicast method is currently on or off on the router, use the **show dhcp relay** command. For information, see the *Behavior for Bound Clients and Address Renewals* section in [“Managing Host Routes Using DHCP Relay Proxy” on page 542](#).

The dhcpRelayGeneral logging event category uses the debug severity level to log DHCP reply packets that are transmitted to clients using a layer 2 unicast address and a layer 3 broadcast address.

The **set dhcp relay broadcast-flag-replies** command configures the router to use the setting of the broadcast flag in DHCP request packets to control the transmission of DHCP reply packets. The **set dhcp relay layer2-unicast-replies** command and the **set dhcp relay broadcast-flag-replies** command are mutually exclusive. For more information, see [“Interaction of DHCP Relay Broadcast Flag with the Layer 2 Unicast Transmission Method” on page 516](#).



NOTE: When you enable the layer 2 unicast transmission feature, the DHCP relay and DHCP relay proxy instance must be the next hop from the DHCP clients. Otherwise, the DHCP reply packets might be discarded.

The layer 2 unicast transmission method is not supported on non-ASIC line modules.

-
- To configure the optional broadcast transmission method:

`host1(config)#set dhcp relay layer2-unicast-replies`

**Related
Documentation**

- [Using the Broadcast Flag Setting to Control Transmission of DHCP Reply Packets on page 516](#)
- [Interaction of DHCP Relay Broadcast Flag with the Layer 2 Unicast Transmission Method on page 516](#)
- `set dhcp relay layer2-unicast-replies`

Using Option 60 Strings to Forward Client Traffic to Specific DHCP Servers

The DHCP functionality supports the DHCP vendor class identifier option (option 60). This support allows DHCP relay to compare option 60 strings in received DHCP client packets against strings that you configure on the router. You can use the DHCP relay option 60 feature when providing converged services in your network environment—option 60 support enables DHCP relay to direct client traffic to the specific DHCP server (the vendor-option server) that provides the service that the client requires. Or, as another option, you can configure option 60 strings to direct traffic to the DHCP local server in the current virtual router.

For example, you might have an environment in which some DHCP clients require only Internet access, while other clients require IPTV service. The clients that need Internet access get their addresses assigned by the DHCP local server on the E Series router (in equal-access mode). Clients requiring IPTV must be relayed to a specific DHCP server that provides the service. To support both types of clients, you configure two option 60 strings on the DHCP relay. Now, when any DHCP client packets are received with option 60 strings configured, the strings are matched against all strings configured on the DHCP relay. If the client string matches the first string you configured, that client is directed to the DHCP local server and gains Internet access. Client traffic with an option 60 string that matches your second string is relayed to the DHCP server that provides the IPTV service. In addition, you can configure a default action, which DHCP relay performs when a client option 60 string does not match any strings you have configured—for example, you might specify that all clients with non-matching strings be dropped.



NOTE: You must enter the **set dhcp relay** command to create and enable DHCP relay before you enter the **set dhcp vendor-option** command to compare option 60 strings in received DHCP client packets against strings that you configure on the router and forward them to specific DHCP servers. Otherwise, the client traffic that matches the configured option 60 strings are dropped.

You use the **set dhcp vendor-option** command to configure vendor-option (option 60) strings to control DHCP client traffic. Create DHCP vendor-option servers by configuring DHCP relay to match DHCP option 60 strings and to specify what action to use for the traffic.

Use the following guidelines when configuring the **set dhcp vendor-option** command:

- Use the **equals** or **starts-with** keywords to specify a unique string to match, and to configure the action to take for traffic with a matching string:
 - **equals**—The DHCP client string is an exact match of the specified string
 - **starts-with**—The DHCP client string is a partial match, from left-to-right, of the specified string. For example, a client string of **day** matches a **starts-with** configured string of **daytime**.
- Use the following keywords to configure actions for matching strings:
 - **local-server**—Forward packets to the DHCP local server
 - **relay**—Forward packets to the DHCP server with the specified IP address
- Use the **default** keyword to set the default action to take when the option 60 string does not match a configured vendor-option string. Use the following keywords to configure actions for nonmatching strings:
 - **drop**—Discard traffic
 - **local-server**—Forward packets to the DHCP local server
 - **proxy-client**—Forward traffic to the DHCP proxy client server
 - **relay**—Forward packets to the DHCP server with the specified IP address
 - **relay-server-list**—Forward traffic to all non-vendor option DHCP servers. The relay-server-list consists of all non-vendor option servers. Non-vendor option servers are those servers that are configured with the **set dhcp relay** command but not with the **set dhcp vendor-option** command.
 - When you configure the first DHCP vendor-option and no default action is specified for a configured DHCP application, the router chooses the default action according to the preference of the DHCP applications. The order of preference from first to last is DHCP local server, DHCP relay, and DHCP proxy client.

You can map multiple strings to the same DHCP server. However, you cannot map the same vendor option string to multiple servers. An error message is displayed in the CLI

interface when you attempt to associate the same option 60 string to more than one server.

You can configure a maximum of 100 option 60 strings per DHCP relay. Strings can contain a maximum of 254 characters.

Client packets that have option 60 configured but have no string specified (a string of 0 length) are treated as nonmatching strings and handled accordingly.

- Related Documentation**
- *set dhcp relay*
 - *set dhcp vendor-option*

Configuring the DHCP Relay Option 60 Attribute for Traffic Forwarding

You use the DHCP relay option 60 feature to specify the action performed on DHCP client traffic. The DHCP relay uses the option 60 string in the client traffic to determine what action to take with the incoming traffic.

The following procedure describes a sample procedure that creates three actions for incoming DHCP client traffic, depending on the traffic's option 60 string.

1. Enable the DHCP relay. Do not specify an IP address when you configure DHCP relay to support vendor-option strings.

```
host1(config)#set dhcp relay
```

2. Configure the action DHCP relay takes when the incoming traffic has an exact option 60 string of myword. DHCP relay forwards this traffic to the DHCP server with an IP address of 192.168.7.7.

```
host1(config)#set dhcp vendor-option equals myword relay 192.168.7.7
```

3. Configure the action DHCP relay takes when the incoming traffic has a partial match, from left-to-right, with an option 60 string you have configured. For this command, matching strings include a, ab, abc, and abcd. DHCP relay forwards matching traffic to the DHCP server with IP address 192.168.15.2.

```
host1(config)#set dhcp vendor-option starts-with abcd relay 192.168.15.2
```

4. Configure the default option 60 action. DHCP relay takes this action when the incoming traffic has an option 60 string that does not match any of the option 60 strings that you have configured. In this example, the traffic is sent to the DHCP local server.

```
host1(config)#set dhcp vendor-option default local-server
```

5. (Optional) View your DHCP relay vendor-option configuration.

```
host1(config)#run show dhcp vendor-option
```

Codes:

- * - the configured vendor-string is an exact-match
- default - all DHCP client packets not matching a configured vendor-string
- implied - the DHCP application is configured but has not been enabled with the vendor-option command
- drop - the DHCP application responsible for the action has not been configured yet therefore all packets for this application will be dropped

Total 3 entries.	
Vendor-option	Action
-----	-----
abcd	relay to 192.168.15.2 (rx: 0)
default(*)	local-server (rx: 0, no-match: 0)
myword(*)	relay to 192.168.7.7 (rx: 0)

- Related Documentation**
- *set dhcp relay*
 - *set dhcp vendor-option*
 - *show dhcp vendor-option*

Relaying DHCP Packets That Originate from a Cable Modem

You can use the DHCP vendor class identifier option (option 60) to configure DHCP relay to relay DHCP packets that originate from a cable modem to an external DHCP server that provides the cable modem with the configuration it requests.

Configure the vendor class identifier option to match the string used by cable modems—DHCP relay then forwards the packets to each DHCP server that you configured with the **set dhcp vendor-option** command (these servers are also considered to be cable-modem DHCP servers).

- To relay DHCP packets from a cable modem:

```
host1(config)#set dhcp relay
host1(config)#service dhcp-local equal-access
host1(config)#set dhcp vendor-option equals docsis relay 192.168.1.1
host1(config)#set dhcp vendor-option equals cablemodem relay 192.168.1.1
```

Use the **show dhcp summary** and **show dhcp vendor-option** commands to display information about the cable modem DHCP relay configuration. See the *Monitoring and Troubleshooting DHCP* chapter in this guide.

- Related Documentation**
- *set dhcp relay*
 - *set dhcp vendor-option*

DHCP Relay Agent Information Option (Option 82) Suboption Values Overview

The DHCP relay agent information option (option 82) enables you to include additional useful information in the client-originated DHCP packets that the DHCP relay forwards to a DHCP server.

When the DHCP relay agent information option is enabled, the DHCP relay adds the option 82 information to packets it receives from clients, then forwards the packets to the DHCP server. The DHCP server uses the option 82 information to decide which IP address to assign to the client—the DHCP server might also use information in the option 82 field for additional purposes, such as determining which services to grant to the client. The DHCP server sends its reply back to the DHCP relay, which removes the option 82 information field from the message, and then forwards the packet to the client.

The option 82 information is made up of a sequence of suboptions. JunosE Software supports the following DHCP relay agent information suboptions.

- Agent Circuit ID (suboption 1)—An ASCII string that identifies the interface on which a client DHCP packet is received.
- Agent Remote ID (suboption 2)—An ASCII string assigned by the relay agent that securely identifies the client.
- Vendor-Specific (suboption 9)—The JunosE Software data field, which contains the Internet Assigned Numbers Authority (IANA) enterprise number (4874) used by JunosE Software and either or both the layer 2 circuit ID and the user packet class.
 - Layer 2 Circuit ID (type 1)—The hexadecimal representation of the layer 2 identifier in the Agent Circuit ID (suboption 1) value (for example, the ATM VPI/VCI or Ethernet SVLAN/VLAN ID.) You can configure this suboption type without the Agent Circuit ID.
 - User Packet Class (type 2)—The hexadecimal representation of the user packet class field, whose value is assigned by the layer 2 policy application. The layer 2 policy application can be used to map the DHCP packet or message IEEE 802.1p value to the user packet class field. See the *JunosE Policy Management Configuration Guide* for information about layer 2 policies.

The Agent Circuit ID suboption (suboption 1) and the Agent Remote ID suboption (suboption 2) are typically determined by the client network access device and depend on the network configuration. The Vendor-Specific suboption (suboption 9) is more flexible and can be used by administrators to associate specific data with the DHCP messages relayed between the DHCP relay and the DHCP server. For example the Vendor-Specific suboption can include the client's IEEE 802.1p value, which identifies the client's user priority.



NOTE: The DHCP relay agent replaces any existing Vendor-Specific value in the client packet with the relay agent's value.

The JunosE Software provides two commands that you can use to configure DHCP relay agent information suboptions.

- The **set dhcp relay agent sub-option** command—Enables you to configure option 82 to include any combination of the supported suboptions, including the Vendor-Specific suboption.
- The **set dhcp relay agent** command—Enables you to configure option 82 to include either or both the Agent Circuit ID suboption (suboption 1) and Agent Remote ID suboption (suboption 2). The command does not support the Vendor-Specific suboption (suboption 9).



NOTE: The `set dhcp relay agent` command is a legacy command, which JunosE Software continues to support to provide backward-compatibility for existing scripts. We recommend that all new configurations use the `dhcp relay agent sub-option` command.

The `set dhcp relay agent sub-option` command enables you to manage specific option 82 suboptions without impacting the configuration of other suboptions. The legacy `set dhcp relay agent` command, however, changes the configuration of suboptions in some cases.

Table 111 on page 527 indicates the effect each command has on enabling or disabling relay agent information suboptions.

Table 111: Effect of Commands on Option 82 Suboption Settings

Command	Suboption and Status		
	Agent Circuit ID	Agent Remote ID	Vendor-Specific
<code>set dhcp relay agent sub-option circuit-id</code>	Enable	No change	No change
<code>set dhcp relay agent sub-option remote-id</code>	No change	Enable	No change
<code>set dhcp relay agent sub-option vendor-specific <i>suboption-type</i></code>	No change	No change	Enable specified suboption type
<code>no set dhcp relay agent sub-option circuit-id</code>	Disable	No change	No change
<code>no set dhcp relay agent sub-option remote-id</code>	No change	Disable	No change
<code>no set dhcp relay agent sub-option vendor-specific <i>suboption-type</i></code>	No change	No change	Disable specified suboption type
<code>set dhcp relay agent</code>	Enable	Enable	Not supported
<code>set dhcp relay agent circuit-id-only</code>	Enable	Disable	Not supported
<code>set dhcp relay agent remote-id-only</code>	Disable	Enable	Not supported
<code>no set dhcp relay agent</code>	Disable	Disable	Disable

Format of the JunosE Data Field in the Vendor-Specific Suboption for Option 82

RFC 4243 describes support for data fields from multiple vendors in the Vendor-Specific suboption for option 82. The JunosE Software DHCP relay agent, however, supports only the JunosE Software data field.

RFC 4243 supports the following format of the Vendor-Specific suboption:

0									1									2									3								
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1				
Code (9)									Length									Enterprise Number 1																	
																		DataLen 1																	
\									Suboption Data 1									\																	

The JunosE Software data field appears after the JunosE Software enterprise number and data length fields in the Vendor-Specific suboption. The format of the JunosE data field is a sequence of type/length/value (TLV) tuples. The type field and length field (the length of the following value field) are each 1 byte in size. The JunosE data length field specifies the total length of all TLV tuples. The JunosE Software enterprise number is 4874 (0x130a.)

The format of the Layer 2 Circuit ID type field (type 1) is hexadecimal. The data field length of a normal non-stacked VLAN is 2 bytes, with the VLAN ID occupying the 12 low-order bits of the value; the 4 high-order bits are 0. The data field length of a stacked VLAN is 4 bytes, with the SVLAN ID occupying the 12 low-order bits of the 2 high-order bytes, and the VLAN ID occupying the 12 low-order bits of the 2 low-order bytes; the unused bits are 0. The data field length of a VPI/VCI is 4 bytes, with the VPI occupying the 8 to 10 low-order bits of the 2 high-order bytes, and the VCI occupying the 16 bits of the 2 low-order bytes; the unused bits are 0.

The format of the UPC data field (type 2) is hexadecimal; its data field length is 1 byte, with the UPC occupying the 4 low-order bits of the value; the 4 high-order bits are 0.

Example 1—The Vendor-Specific suboption for a VLAN ID of 2468 (0x09a4) and a UPC of 5 is formatted as follows:

```
09 0c 00 00 13 0a 07 01 02 09 a4 02 01 05
| | |      | | | |   | | |
| | |      | | | |   | |   UPC val: 5
| | |      | | | |   |     UPC len: 1 byte
| | |      | | | |       UPC type: 2
| | |      | | | |   L2 Circuit ID val: 09 a4
| | |      | | |   L2 Circuit ID len: 2 bytes
| | |      |   L2 Circuit ID type: 1
| | |        JUNOSE data len: 7 bytes
| |   JUNOSE IANA: 13 0a
| subopt 9 len: 12 bytes
subopt code: 9
```

Example 2—The Vendor-Specific suboption for a VLAN ID of 135-2468 (0x87-0x09a4, format <SVLAN ID>-<VLAN ID>) and a UPC of 5 is formatted as follows:

```
09 0e 00 00 13 0a 09 01 04 00 87 09 a4 02 01 05
|   |   |           |   |   |   |           |   |   |
|   |   |           |   |   |   |           |   |   | UPC val: 5
```



```
| | | | | | | UPC len: 1 byte
| | | | | | | UPC type: 2
| | | | | | L2 Circuit ID val: 00 87 09 a4
| | | | | | L2 Circuit ID len: 4 bytes
| | | | | | L2 Circuit ID type: 1
| | | | | | JUNOSE data len: 9 bytes
| | JUNOSE IANA: 13 0a
| subopt 9 len: 14 bytes
subopt code: 9
```

Example 3—The Vendor-Specific suboption for a VPI/VCI of 123.45678 (0x7b.0xb26e, format <VPI>.<VCI>) and a UPC of 5 is formatted as follows:

```
09 0e 00 00 13 0a 09 01 04 00 7b b2 6e 02 01 05
| | | | | | | | | | | | | | |
| | | | | | | | | | | UPC val: 5
| | | | | | | | | | | UPC len: 1 byte
| | | | | | | | | | | UPC type: 2
| | | | | | | L2 Circuit ID val: 00 7b b2 6e
| | | | | | | L2 Circuit ID len: 4 bytes
| | | | | | | L2 Circuit ID type: 1
| | | | | | | JUNOSE data len: 9 bytes
| | JUNOSE IANA: 13 0a
| subopt 9 len: 14 bytes
subopt code: 9
```

Related Documentation

- [Configuring the DHCP Relay Agent Option 82 Settings on page 529](#)

Configuring the DHCP Relay Agent Option 82 Settings

You can configure the option 82 attribute for DHCP relay using the following set of tasks

- [Configuring Relay Agent Option 82 Information on page 529](#)
- [Preventing Option 82 Information from Being Stripped from Trusted Client Packets on page 530](#)
- [Configuring Relay Agent Information Option \(Option 82\) Suboption Values on page 530](#)

Configuring Relay Agent Option 82 Information

You can specify the type the relay agent option 82 information that the router adds to DHCP packets before it relays the packets to the DHCP server. You can use one of the following keywords to add either the hostname or virtual router name to the front of the Circuit-Id field or to strip the subinterface ID from the Interface-Id field:

- **hostname**—Adds the router's hostname to the front of the Circuit-Id field; a colon separates the hostname from the circuit information
- **vrname**—Adds the router's virtual router name to the front of the Circuit-Id field; a colon separates the virtual router name from the circuit information
- Use the **exclude-subinterface-id** to strip the subinterface ID from the Interface-Id field. When the interface ID is constructed, it contains the slot/port numbers, the subinterface ID, and the VPI/VCI for ATM interfaces or the VLAN ID for Ethernet interfaces. Use this keyword to remove the subinterface ID from the Interface-Id field.

The **hostname** and **vrname** keywords are a toggle; that is, specifying either hostname or virtual router name turns off the other selection.

- To configure the relay agent option 82 information:

```
host1(config)#set dhcp relay options hostname
```

Preventing Option 82 Information from Being Stripped from Trusted Client Packets

You can configure DHCP relay or DHCP relay proxy to preserve option 82 information for trusted clients. This ensures that DHCP relay and DHCP relay proxy prevent option 82 information from being stripped off packets destined for a trusted client. A trusted client has a giaddr value of 0. If DHCP relay is configured not to remove option 82 and the giaddr field is 0, option 82 information remains in the packets.

- To prevent the option 82 information from being removed from packets destined for a trusted client:

```
host1(config)#set dhcp relay preserve-trusted-client-option
```

Configuring Relay Agent Information Option (Option 82) Suboption Values

The JunosE Software provides two commands that you can use to configure DHCP relay agent information suboptions.

- To configure option 82 to include any combination of the supported suboptions, including the Vendor-Specific suboption:

```
host1(config)#set dhcp relay agent sub-option remote-id
host1(config)#set dhcp relay agent sub-option circuit-id
host1(config)#set dhcp relay agent sub-option vendor-specific layer2-circuit-id
host1(config)#set dhcp relay agent sub-option vendor-specific user-packet-class
```

- To configure option 82 to include either or both the Agent Circuit ID suboption (suboption 1) and Agent Remote ID suboption (suboption 2). The command does not support the Vendor-Specific suboption (suboption 9).

```
host1(config)#set dhcp relay agent
```

Related Documentation

- [DHCP Relay Agent Information Option \(Option 82\) Suboption Values Overview on page 525](#)
- *set dhcp relay*
- *set dhcp relay agent sub-option*
- *set dhcp relay options*
- *set dhcp relay preserve-trusted-client-option*

Using the set dhcp relay agent sub-option Command to Enable Option 82 Suboption Support



NOTE: We recommend that you use the **set dhcp relay agent sub-option** command for new option 82 suboption configurations. However, JunosE Software continues to support the **set dhcp relay agent** command, with option 82 suboptions, to provide backward-compatibility for existing scripts.

You use the **set dhcp relay agent sub-option** command to enable support for a specific DHCP relay agent option 82 suboption—Agent Circuit ID (suboption 1), Agent Remote ID (suboption 2), and Vendor-Specific (suboption 9). When you issue this command, the router adds DHCP relay agent information suboption 1 to every packet it relays from a DHCP client to a DHCP server. The Agent Circuit ID suboption identifies the interface on which DHCP packets are received. When the packets are received on a LAG interface, the router clearly identifies the interface.

The suboptions include information from the DHCP relay agent that the DHCP server can use to implement parameter assignment policies. The DHCP server echoes the suboptions when it replies to the DHCP client, but the DHCP relay strips the suboptions before relaying the packets to the client.

The Agent Circuit ID suboption identifies the interface on which the DHCP packets are received. This suboption contains the following information, based on interface type:

- ATM interface

```
[<hostname>|<vrname>:]<interface type> <slot>/<port>[.<sub-if>]:<vpi>.<vci>
```

Examples:

```
atm 4/1.2:0.101
relayVr:atm 4/1:0.101
bostonHost:atm 4/1.2:0.101
```

- Ethernet interface

```
[<hostname>|<vrname>:]<interface type> <slot>/<port>
```

Examples:

```
fastEthernet 1/2
relayVr:fastEthernet 1/2
bostonHost:fastEthernet 1/2
```

- Ethernet interface with VLAN

```
[<hostname>|<vrname>:]<interface type> <slot>/<port>[.<sub-if>]:<vlan id>
```

Examples:

```
fastEthernet 1/2.3:4
relayVr:fastEthernet 1/2:4
bostonHost:fastEthernet 1/2.3:4
```

- Ethernet interface with Stacked VLAN

```
[<hostname>|<vrname>:]<interface type> <slot>/<port>[.<sub-if>]:  
  <svlan id>-<vlan id>
```

Examples:

```
fastEthernet 1/2.3:4-5  
relayVr:fastEthernet 1/2:4-5  
bostonHost:fastEthernet 1/2.3:4-5
```

- LAG interface

```
[<hostname>|<vrname>:]<interface type> <bundle name>
```

Examples:

```
lag bundleA  
relayVr:lag bundleA  
bostonHost:lag bundleA
```

- LAG interface with VLAN

```
[<hostname>|<vrname>:]<interface type> <bundle name>[.<sub-if>]:<vlan id>
```

Examples:

```
lag bundleA.1:2  
relayVr:lag bundleA:2  
bostonHost:lag bundleA.1:2
```

- LAG interface with Stacked VLAN

```
[<hostname>|<vrname>:]<interface type> <bundle name>[.<sub-if>]:  
  <svlan id>-<vlan id>
```

Examples:

```
lag bundleA.1:2-3  
relayVr:lag bundleA:2-3  
bostonHost:lag bundleA.1:2-3
```

The Agent Remote ID suboption contains a value only when (1) the interface is a dynamic ATM interface and (2) the **subscriber** command is used to configure a username and domain name for the interface. If both conditions are met, the suboption contains a string with the username and domain name in the format: *username@domainname*.

The Vendor-Specific suboption contains a value that includes a JunosE data field. You can configure the data field to support one or both of the following values:

- **layer2-circuit-id** (type 1)—The hexadecimal representation of the layer 2 identifier in the Agent Circuit ID (suboption 1) value (for example, the ATM VPI/VCI or Ethernet SVLAN/VLAN ID). You can configure this suboption type without the Agent Circuit ID.
- **user-packet-class** (type 2)—The hexadecimal representation of the user packet class field, whose value is assigned by the layer 2 policy application. The layer 2 policy application can be used to map the DHCP packet or message IEEE 802.1p value to the user packet class field. See the *JunosE Policy Management Configuration Guide* for information about layer 2 policies.

- Related Documentation**
- [DHCP Relay Agent Information Option \(Option 82\) Suboption Values Overview on page 525](#)
 - [Configuring the DHCP Relay Agent Option 82 Settings on page 529](#)
 - *set dhcp relay agent sub-option*

Using the set dhcp relay agent Command to Enable Option 82 Suboption Support



NOTE: The **set dhcp relay agent** command, when used to configure option 82 suboptions is a legacy command, which JunosE Software continues to support to provide backward-compatibility for existing scripts. We recommend that you use the **dhcp relay agent sub-option** command for new option 82 suboption configurations.

You can use the **set dhcp relay agent** command to enable support for DHCP relay agent option, which includes the option 82 suboptions—Agent Circuit ID (suboption 1) and Agent Remote ID (suboption 2). This command does not support the Vendor-Specific option (suboption 9).

The suboptions include information from the DHCP relay agent that the DHCP server can use to implement parameter assignment policies. The DHCP server echoes the suboptions when it replies to the client—the DHCP relay agent can optionally strip the option 82 information before relaying the packets to the client. (Use the CLI command **set dhcp relay preserve-trusted-client-option** to configure this behavior for trusted clients.)

When you issue the **set dhcp relay agent** command, the router adds the configured DHCP relay agent information suboptions to every packet it relays from a DHCP client to a DHCP server.

The **circuit-id-only** keyword specifies the Agent Circuit ID suboption, which contains the following information, based on interface type. This keyword disables support for the Agent Remote ID suboption.

- ATM interface

```
[<hostname>|<vname>:]<interface type> <slot>/<port>[.<sub-if>]:<vpi>.<vci>
```

Examples:

```
atm 4/1.2:0.101
relayVr:atm 4/1:0.101
bostonHost:atm 4/1.2:0.101
```

- Ethernet interface

```
[<hostname>|<vname>:]<interface type> <slot>/<port>
```

Examples:

```
fastEthernet 1/2
relayVr:fastEthernet 1/2
```

```
bostonHost:fastEthernet 1/2
```

- Ethernet interface with VLAN

```
[<hostname>|<vrname>:]<interface type> <slot>/<port>[.<sub-if>]:<vlan id>
```

Examples:

```
fastEthernet 1/2.3:4
relayVr:fastEthernet 1/2:4
bostonHost:fastEthernet 1/2.3:4
```

- Ethernet interface with Stacked VLAN

```
[<hostname>|<vrname>:]<interface type> <slot>/<port>[.<sub-if>]:
<svlan id>-<vlan id>
```

Examples:

```
fastEthernet 1/2.3:4-5
relayVr:fastEthernet 1/2:4-5
bostonHost:fastEthernet 1/2.3:4-5
```

- LAG interface

```
[<hostname>|<vrname>:]<interface type> <bundle name>
```

Examples:

```
lag bundleA
relayVr:lag bundleA
bostonHost:lag bundleA
```

- LAG interface with VLAN

```
[<hostname>|<vrname>:]<interface type> <bundle name>[.<sub-if>]:<vlan id>
```

Examples:

```
lag bundleA.1:2
relayVr:lag bundleA:2
bostonHost:lag bundleA.1:2
```

- LAG interface with Stacked VLAN

```
[<hostname>|<vrname>:]<interface type> <bundle name>[.<sub-if>]:
<svlan id>-<vlan id>
```

Examples:

```
lag bundleA.1:2-3
relayVr:lag bundleA:2-3
bostonHost:lag bundleA.1:2-3
```

The **remote-id-only** keyword specifies the Agent Remote ID suboption, which contains a value only when (1) the interface is a dynamic ATM interface and (2) the **subscriber** command is used to configure a username and domain name for the interface. If both conditions are met, the suboption contains a string with the username and domain name in the format: *username@domainname*. The **remote-id-only** keyword disables support for the Agent Circuit ID suboption.

If you do not explicitly specify the **circuit-id-only** or **remote-id-only** keyword, both suboptions are used.

Related Documentation

- [DHCP Relay Agent Information Option \(Option 82\) Suboption Values Overview on page 525](#)
- [Configuring the DHCP Relay Agent Option 82 Settings on page 529](#)
- `set dhcp relay`

Example: Using DHCP Relay Option 82 to Pass IEEE 802.1p Values to DHCP Servers

This example shows how you can use the DHCP relay agent option 82 feature to enable the DHCP server to assign an IP address to provide the desired service to the DHCP client.

- [Requirements on page 535](#)
- [Overview on page 535](#)
- [Configuring DHCP Relay Option 82 to Pass IEEE 802.1p Values to DHCP Servers on page 537](#)

Requirements

This example uses the following software and hardware components:

- JunosE Release 7.1.0 or higher-numbered releases
- E Series router (ERX7xx models, ERX14xx models, the ERX310 router, the E120 router, or the E320 router)
- ASIC-based line modules that support Fast Ethernet or Gigabit Ethernet

Overview

Using the DHCP relay agent option 82 feature, you can configure an environment in which a customized DHCP server assigns an IP address that provides the desired service to the DHCP client.

The DHCP server uses information based on the IEEE 802.1p values, which are extracted from the DHCP packets using JunosE Software layer 2 policies, to determine the appropriate IP address to assign to the client.

This type of environment, which is illustrated in [Figure 13 on page 536](#), includes the following components:


- Layer 2 policy on the ingress interface (that is, the interface that receives the client's DHCP packet) that maps the 802.1p value from the packet to a user packet class (UPC.)



NOTE: To ensure optimal performance when mapping 802.1p values to UPCs, order the classifier groups in the VLAN policy list with the most often used UPC values listed first.

- DHCP relay agent option 82 configuration that enables Vendor-Specific suboption type 2 (User Packet Class) support and maps the Layer 2 policy user packet class to the option 82 user packet class suboption.
- Customized DHCP server configuration that assigns IP addresses based on the option 82 user packet class suboption. The IP address is associated with the appropriate quality, type, or class of service for the user packet class specified in the option 82 suboption.

Figure 13: Passing 802.1p Values to the DHCP Server



Configuring DHCP Relay Option 82 to Pass IEEE 802.1p Values to DHCP Servers

Step-by-Step Procedure The following example describes a sample procedure that creates an environment that passes 802.1p values to the DHCP server, which then assigns an IP address that enables the desired service to the DHCP client:

1. Configure a layer 2 policy that maps 802.1p values to user packet class values for a VLAN interface.

```
host1(config)# vlan classifier-list dot1p0 user-priority 0
host1(config)# vlan classifier-list dot1p1 user-priority 1
host1(config)# vlan classifier-list dot1p2 user-priority 2
host1(config)# vlan classifier-list dot1p3 user-priority 3
host1(config)# vlan classifier-list dot1p4 user-priority 4
host1(config)# vlan classifier-list dot1p5 user-priority 5
host1(config)# vlan classifier-list dot1p6 user-priority 6
host1(config)# vlan classifier-list dot1p7 user-priority 7
host1(config)# vlan policy-list dot1pToUpc
host1(config-policy-list)# classifier-group dot1p0
host1(config-policy-list-classifier-group)# user-packet-class 0
host1(config-policy-list-classifier-group)#exit
host1(config-policy-list)# classifier-group dot1p1
host1(config-policy-list-classifier-group)# user-packet-class 1
host1(config-policy-list-classifier-group)#exit
host1(config-policy-list)# classifier-group dot1p2
host1(config-policy-list-classifier-group)# user-packet-class 2
host1(config-policy-list-classifier-group)#exit
host1(config-policy-list)# classifier-group dot1p3
host1(config-policy-list-classifier-group)# user-packet-class 3
host1(config-policy-list-classifier-group)#exit
host1(config-policy-list)# classifier-group dot1p4
host1(config-policy-list-classifier-group)# user-packet-class 4
host1(config-policy-list-classifier-group)#exit
host1(config-policy-list)# classifier-group dot1p5
host1(config-policy-list-classifier-group)# user-packet-class 5
host1(config-policy-list-classifier-group)#exit
host1(config-policy-list)# classifier-group dot1p6
host1(config-policy-list-classifier-group)# user-packet-class 6
host1(config-policy-list-classifier-group)#exit
host1(config-policy-list)# classifier-group dot1p7
host1(config-policy-list-classifier-group)# user-packet-class 7
host1(config-policy-list-classifier-group)#exit
host1(config-policy-list)#exit
host1(config)# profile atm1483BaseProfile
host1(config-profile)# vlan policy input dot1pToUpc statistics enabled
host1(config-profile)#exit
host1(config)#
```

2. (Optional) Verify the policy list configuration.

```
host1(config)# run show policy-list dot1pToUpc
```

Policy Table

```
VLAN Policy dot1pToUpc
Administrative state: enable
Reference count:      1
Classifier control list: dot1p0, precedence 100
                        user-packet-class 0
Classifier control list: dot1p1, precedence 100
```

```

user-packet-class 1
Classifier control list: dot1p2, precedence 100
user-packet-class 2
Classifier control list: dot1p3, precedence 100
user-packet-class 3
Classifier control list: dot1p4, precedence 100
user-packet-class 4
Classifier control list: dot1p5, precedence 100
user-packet-class 5
Classifier control list: dot1p6, precedence 100
user-packet-class 6
Classifier control list: dot1p7, precedence 100
user-packet-class 7

Referenced by interface(s):
None

Referenced by profile(s):
atm1483BaseProfile input policy, statistics enabled

Referenced by merged policies:
None

```

3. Configure the DHCP relay to use the option 82 suboptions. This configuration includes the command that specifies the mapping of the user packet class values from the layer 2 policy to the user-packet-class type in the option 82 Vendor-Specific suboption.

```

host1(config)# set dhcp relay 192.168.32.1 proxy
host1(config)# set dhcp relay 192.168.32.2
host1(config)# set dhcp relay agent sub-option circuit-id
host1(config)# set dhcp relay agent sub-option remote-id
host1(config)# set dhcp relay agent sub-option vendor-specific
user-packet-class
host1(config)# set dhcp relay agent sub-option vendor-specific
layer2-circuit-id
host1(config)# set dhcp relay options hostname
host1(config)# set dhcp relay options exclude-subinterface-id
host1(config)# set dhcp relay inhibit-access-route-creation
host1(config)# set dhcp relay trust-all
host1(config)# set dhcp relay override agent-option

```

4. (Optional) Verify the DHCP Relay configuration.

```

host1(config)# run show dhcp relay

DHCP Relay Configuration
-----
Mode: Proxy
  Restore Client Timeout: 72
  Inhibit Access Route Creation: off
  Assign Giaddr to Source IP: off
  Layer 2 Unicast Replies: off
  Giaddr Selects Interface: off
  Relay Agent Information Option (82):
    Override Giaddr: off
    Override Option: on
    Trust All Clients: on
    Preserve Option From Trusted Clients: off
    Circuit-ID Sub-option (1): on
      select - hostname
      select - exclude-subinterface-id
    Remote-ID Sub-option (2): on

```

```

Vendor-Specific Sub-option (9): on
  select - layer2-circuit-id
  select - user-packet-class

DHCP Server Addresses
-----
192.168.32.1
192.168.32.2

```

Related Documentation

- [Configuring the DHCP Relay Agent Option 82 Settings on page 529](#)
- [Configuring Relay Agent Option 82 Information on page 529](#)
- *classifier-group*
- *set dhcp relay*
- *set dhcp relay agent sub-option*
- *set dhcp relay assign-giaddr-source-ip*
- *set dhcp relay broadcast-flag-replies*
- *set dhcp relay giaddr-selects-interface*
- *set dhcp relay layer2-unicast-replies*
- *set dhcp relay options*
- *set dhcp relay override*
- *set dhcp relay preserve-trusted-client-option*
- *set dhcp relay trust-all*
- *set dhcp vendor-option*
- *vlan classifier-list*

Rate of DHCP Client Packets Processed by DHCP Relay Overview

In instances when multiple IP clients send a large number of packets, DHCP relay is kept busy processing client packets. While DHCP relay processes client packets, it does not process the packets returned from DHCP server at the same rate. This causes a drop in the number of successful transactions between the remote clients and DHCP server.

You can avoid this situation by limiting the number of client packets that DHCP relay processes and assigning a higher priority to packets that DHCP server sends. While DHCP relay monitors the load on the uplink line module, it calculates the number of client packets it can process for the next second. This calculation enables DHCP relay to handle only that number of client packets and process all the packets as calculated. DHCP relay continues to monitor the uplink line module and sets a new rate for packets processed every second. DHCP assigns higher priority to server packets so, when processing an excessive number of packets, DHCP relay discards the client packets first.

Manually Configuring the Maximum Rate of Client Packets Processed Per Second by DHCP Relay

DHCP relay monitors the load on the interface controller (IC). In some cases, when packets are discarded before they reach the IC, the DHCP relay cannot set the maximum client packet rate automatically. You can then manually reconfigure the rate at which DHCP relay processes client packets per second. Manual reconfiguration of the maximum client packet rate is necessary in cases such as the following:

- When DHCP relay forwards a large number of client packets, DHCP server might not be able to process all of them. DHCP server discards the excess packets. You can set the maximum rate of client packets based on the server capability.
- When the uplink line module cannot handle heavy loads, packets are discarded before they reach the IC. You can set the maximum rate of client packets based on the uplink load capacity.
- When DoS parameters are configured on the uplink line module, packets are discarded at the forwarding controller (FC). You can set the maximum rate of client packets based on these DoS parameters.

Related Documentation

- [Configuring the Rate of Client Packets Processed by DHCP Relay on page 540](#)

Configuring the Rate of Client Packets Processed by DHCP Relay

You can control the maximum number of client packets that DHCP relay forwards to the DHCP server.

To limit the maximum number of client packets that DHCP relay forwards:

- From Global Configuration mode, set the maximum number of client packets that DHCP relay handles.

```
host1(config)#set dhcp relay max-client-packet-rate 1024
```

Related Documentation

- [Rate of DHCP Client Packets Processed by DHCP Relay Overview on page 539](#)
- `set dhcp relay max-client-packet-rate`

Configuring DHCP Relay Proxy

The DHCP relay proxy is an enhancement to the E Series router's DHCP relay component. The DHCP relay proxy manages host routes for DHCP clients, and determines which offer to use when there are multiple DHCP servers configured.



NOTE: The E Series router configured as a DHCP relay proxy must be the first hop from the DHCP client. If it is not the first hop, the router defaults to the DHCP relay configuration.

You can configure a DHCP relay proxy using the following set of tasks:

- [Enabling DHCP Relay Proxy on page 541](#)
- [Use the First Offer from a DHCP Server on page 541](#)
- [Set a Timeout for DHCP Client Renewal Messages on page 541](#)

Enabling DHCP Relay Proxy

Enable DHCP relay proxy and specify an IP address for the DHCP server. After you are in DHCP relay proxy mode, all **set dhcp relay** commands are supported.

```
host1(config)#set dhcp relay 192.168.29.10 proxy
```

When you issue this command, the router adds the IP address to the list of DHCP servers (up to five) and forwards all request packets to all configured servers.

After you are in DHCP relay proxy mode, all **set dhcp relay** commands are supported.

Use the First Offer from a DHCP Server

You can configure the DHCP relay proxy to use the first offer it receives from any configured DHCP server and send that offer to the DHCP client. By default, DHCP relay proxy sends the most appropriate offer it receives from the configured DHCP servers to the DHCP client.

```
host1(config)#set dhcp relay proxy send-first-offer
```

Set a Timeout for DHCP Client Renewal Messages

You can set the amount of time, in the range 1–168 hours, that the DHCP relay proxy waits for a renewal message from DHCP clients after a router reboot or switchover occurs. A renewal message is required from DHCP clients when a router reboot or switchover occurs. If no renewal message is received before the timeout expires, the relay proxy declares the client no longer active and removes the client's host route. By default, DHCP relay proxy uses timeout of 72 hours.

```
host1(config)#set dhcp relay proxy timeout 8
```



NOTE: DHCP relay proxy does not remove a DHCP client's host route when the lease for the client's IP address expires. DHCP relay proxy will instead remove the host route when the relay proxy timeout expires. To prevent a host route from remaining long after lease expiration, modify the relay proxy timeout from its default setting of 72 hours to a setting close to, but not less than the lease time.

Related Documentation

- *set dhcp relay*
- *set dhcp relay proxy send-first-offer*
- *set dhcp relay proxy timeout*

Managing Host Routes Using DHCP Relay Proxy

The DHCP relay proxy feature enables the E Series router to efficiently manage host routes for DHCP clients, including:

- Installing routes when DHCP clients are configured
- Removing routes when DHCP clients release their DHCP-assigned addresses or when the addresses expire

When a DHCP client sends a request to an external DHCP server, the relay proxy receives the request and forwards it to the external DHCP server. The relay proxy then sends the DHCP server's response back to the client. This process is similar to that used by the DHCP relay component. The DHCP client views the relay proxy as a DHCP server, and the DHCP server sees the relay proxy as a DHCP relay agent.

To DHCP clients, there is no difference when they use a DHCP relay or a DHCP relay proxy. However, the DHCP relay proxy differs from the DHCP relay in how client address renewals and releases are handled:

- With the DHCP relay proxy, DHCP clients communicate with the relay proxy to renew and release addresses.
- With the DHCP relay, DHCP clients communicate directly with the DHCP server to renew and release addresses.

A major benefit of the relay proxy configuration is that the E Series router is kept informed of the status of a DHCP client's address. When addresses are released by clients, the router removes the installed host route for that client. In the DHCP relay configuration, the router does not know when addresses have been renewed or released; the host routes that are no longer needed are still unavailable.

For additional information on managing client bindings, see [“Viewing and Deleting DHCP Client Bindings” on page 466](#).

Selecting the DHCP Server Response

Similar to the DHCP relay, the DHCP relay proxy enables you to specify up to five DHCP servers to provide address and configuration information for a DHCP client. As an added benefit over the relay, when using multiple DHCP external servers, you can configure how the DHCP relay proxy determines which offer to send to the DHCP client. You can configure the DHCP relay proxy to use either the single best offer or the first offer it receives from the DHCP servers.

If there are multiple offers, the DHCP relay proxy selects the final offer based on the following priorities:

1. The offer that contains the IP address requested by the DHCP client.
2. The offer that contains an IP address on the same subnetwork as the requested IP address.
3. The offer that has the longest lease time.

If you have enabled the optional select-first-offer feature, the DHCP relay proxy immediately uses the first offer that it receives from any DHCP server.

Behavior for Bound Clients and Address Renewals

When a DHCP client is already bound to an IP address or is renewing the lease on its IP address, DHCP relay proxy unicasts DHCP ACK and DHCP NAK replies to the client regardless of the current configuration of the **set dhcp relay layer2-unicast-replies** command or the **set dhcp relay broadcast-flag-replies** command. These commands control the transmission method used for DHCP reply packets.

This behavior applies only to DHCP relay proxy; it does not apply to DHCP relay because DHCP relay does not maintain a list of active clients or receive address renewal requests from clients.

For information about using the **set dhcp relay layer2-unicast-replies** command, see [“Configuring the Layer 2 Unicast Transmission Method for Reply Packets to DHCP Clients” on page 521](#). For information about using the **set dhcp relay broadcast-flag-replies** command, see [“Configuring the Layer 2 Unicast Transmission Method for Reply Packets to DHCP Clients” on page 521](#).

Related Documentation

- [Interaction of DHCP Relay Broadcast Flag with the Layer 2 Unicast Transmission Method on page 516](#)
- [Using the Broadcast Flag Setting to Control Transmission of DHCP Reply Packets on page 516](#)
- *set dhcp relay*
- *set dhcp relay broadcast-flag-replies*
- *set dhcp relay layer2-unicast-replies*

CHAPTER 22

Configuring the DHCP External Server Application

The following sections describe how to configure the DHCP external server application on the E Series router:

- [DHCP External Server Overview on page 545](#)
- [Preservation of Dynamic Subscriber Interfaces with DHCP External Server Overview on page 547](#)
- [DHCP External Server Identification of Clients with Duplicate MAC Addresses Overview on page 548](#)
- [DHCP External Server Configuration Requirements on page 550](#)
- [Enabling and Disabling the DHCP External Server Application on page 550](#)
- [Monitoring DHCP Traffic Between Remote Clients and DHCP Servers on page 551](#)
- [Synchronizing the DHCP External Application and the Router on page 551](#)
- [Configuring Interoperation with Ethernet DSLAMs on page 551](#)
- [Configuring the DHCP External Server to Support the Creation of Dynamic Subscriber Interfaces on page 552](#)
- [Configuring DHCP External Server to Control Preservation of Dynamic Subscriber Interfaces on page 553](#)
- [Configuring Dynamic Subscriber Interfaces for Interoperation with DHCP Relay and DHCP Relay Proxy on page 554](#)
- [Deleting Clients from a Virtual Router's DHCP Binding Table on page 556](#)
- [Configuring DHCP External Server to Uniquely Identify Clients with Duplicate MAC Addresses on page 557](#)
- [Configuring DHCP External Server to Re-Authenticate Auto-Detected Dynamic Subscriber Interfaces on page 558](#)

DHCP External Server Overview

You can configure the E Series router to provide support for an external DHCP server. This enables the router, which is not running DHCP relay or DHCP proxy server, to monitor DHCP packets and to keep information for subscribers based on their IP address and

MAC address. When the E Series router's DHCP external server application is used, all DHCP traffic to and from the DHCP server is monitored by the router.


The services provided by integrating the E Series router's DHCP external server application with SRC software are similar to those provided when the DHCP local server is integrated with SRC software. The router's DHCP external application is used together with other features of the router to provide subscriber management. For additional information, see the *Configuring Subscriber Management* chapter.



NOTE: To ensure that DHCP external server with DHCP relay proxy processes unicast reply packets (such as renewal ACK and NAK packets), you must configure DHCP external server with the IP address of the DHCP relay proxy's giaddr. This configuration ensures that DHCP external server processes renewal ACK packets, which in turn enables the updating of client leases.

Figure 14 on page 546 shows a network that includes an external DHCP server and the E Series router.

Figure 14: DHCP External Server



In Figure 14 on page 546, the subscriber requests an address from the DHCP server through the E Series router. All communication between the subscriber and the DHCP server is monitored by the E Series router. After the subscriber receives an IP address, the subscriber is able to access the Internet and use the value-added services provided by the E Series router and by the SAE software. For this to occur, the edge network must be using a DHCP relay function.

When the subscriber requests an IP address from the DHCP server, the E Series router performs the following actions:

- Identifies the subscriber's IP address, MAC address, giaddr, and client identifier
- Extracts the lease time, creates a shadow lease, and starts its own lease timer that is associated with the subscriber

The E Series router views the subscriber as active once the subscriber sends a packet. The router then performs the following actions:

- Processes the subscriber's IP address by using a route map
- Extracts the dynamic subscriber interface profile (optional)

- Creates the subscriber's dynamic subscriber interface

If the SRC software is configured, the router also performs the following actions:

- Alerts the SRC software that the dynamic subscriber interface exists
- Alerts the SRC software that the subscriber's address exists and provides DHCP options

The SRC software then provides its enhanced services to the subscriber.

The E Series router monitors all traffic between the subscriber and the DHCP server, and resets the shadow lease by monitoring the DHCP server/client lease renewal. When the subscriber disconnects, the shadow lease will eventually expire. The E Series router then performs the following actions:

- Deletes the subscriber's dynamic subscriber interface
- Alerts the SRC software that the dynamic subscriber interface has been deleted
- Alerts the SRC software that the subscriber's address has been deleted

For additional information on managing client bindings, see [“Viewing and Deleting DHCP Client Bindings”](#) on page 466.

Preservation of Dynamic Subscriber Interfaces with DHCP External Server Overview

You can configure the DHCP external server application to control whether the router preserves or deletes and re-creates a DHCP client's existing dynamic subscriber interface in certain situations.

The DHCP discovery process assigns an IP address to a DHCP client. A client initiates the discovery process on a primary IP interface in the router. When this process completes successfully, the IP subscriber manager application may create a dynamic subscriber interface for the client that exists with the client's primary interface. A client normally receives broadcast traffic, such as the traffic associated with the DHCP discovery process, on its primary interface. A client normally receives unicast traffic, such as the traffic associated with the DHCP renewal process, on its dynamic subscriber interface if one exists.

A DHCP client that has successfully completed the discovery process and has been assigned an IP address in the DHCP external server application is referred to as a *bound client*. An IP address is leased to a client for a specified period of time. Before the lease period expires, most bound DHCP clients typically use the DHCP renewal process to extend their IP address lease. However, some bound DHCP clients might extend their IP address lease by restarting the DHCP discovery process instead of using the DHCP renewal process.

When a bound DHCP client on a dynamic subscriber interface extends its address lease by restarting the discovery process on its primary IP interface, you can configure the DHCP external server application to control whether the client's existing dynamic subscriber interface is preserved, or deleted and re-created. By default, the DHCP external server preserves the client's existing dynamic subscriber interface in this situation. To configure the DHCP external server to delete and re-create the client's dynamic subscriber interface

after the client restarts discovery, you must issue the **ip dhcp-external recreate-subscriber-interface** command from Global Configuration mode.

When a bound DHCP client restarts the discovery process on a different primary IP interface than the interface on which it initiated the original discovery process, the DHCP external server application always deletes and re-creates the existing dynamic subscriber interfaces for that client.

You must use the **ip dhcp-external recreate-subscriber-interface** command within a specific virtual router context. Because you issue this command on a per-virtual router basis, different virtual routers configured in the same router can use different settings for this command.

In some lower-numbered JunosE Software releases, the default behavior for the DHCP external server was to delete and re-create the dynamic subscriber interface after a bound client restarted the discovery process on its primary IP interface. If you are upgrading the JunosE Software on the router from one of these releases to the current release, keep in mind that you must explicitly issue the **ip dhcp-external recreate-subscriber-interface** command to continue to delete and re-create the dynamic subscriber interface. The router no longer deletes and re-creates the dynamic subscriber interface by default in this situation.

See *DHCP External Server* in the *Known Behavior* section of the *JunosE Release Notes* for a list of the JunosE releases in which deleting and re-creating the dynamic subscriber interface was the default behavior for the DHCP external server.

**Related
Documentation**

- [Configuring DHCP External Server to Control Preservation of Dynamic Subscriber Interfaces on page 553](#)
- *ip dhcp-external recreate-subscriber-interface*

DHCP External Server Identification of Clients with Duplicate MAC Addresses Overview

You can configure the DHCP external server application to use a combination of the media access control (MAC) address and the gateway IP address (giaddr) to uniquely identify DHCP clients attached to the router. Using this feature enables you to manage DHCP clients in network environments in which MAC addresses are not unique.

In some network environments where the DHCP external server application manages DHCP clients from multiple DHCP relays, the same MAC address might be assigned to more than one DHCP client. This can occur, for example, when network adapters are manufactured with the same hardware address, resulting in duplicate MAC addresses among the DHCP clients attached to the router.

To better manage DHCP clients in network environments with multiple DHCP relays in which MAC addresses are not unique, you can configure the DHCP external server application to use a combination of the MAC address and the giaddr to uniquely identify the clients connected to the router. This setting for DHCP external server is also referred to as *duplicate MAC mode*.

By default, DHCP external server uses only the MAC address to uniquely identify DHCP clients. The default setting for DHCP external server is also referred to as *unique MAC mode*.

To enable duplicate MAC mode for the DHCP external server application, you must issue the **dhcp-external duplicate-mac-address** command from Global Configuration mode. To restore the default behavior and re-enable unique MAC mode, issue the **no dhcp-external duplicate-mac-address** command.

Configuration Guidelines for Using Duplicate MAC Mode

Observe the following guidelines when you configure the DHCP external server application to use a combination of the MAC address and giaddr to uniquely identify DHCP clients, otherwise known as enabling duplicate MAC mode:

- Unlike other commands for configuring DHCP external server, the **dhcp-external duplicate-mac-address** command applies globally to all instances of the DHCP external server application on the router, and is not issued on a per-VR basis.
- Although the same MAC address can be assigned to more than one DHCP client in the network, MAC addresses must be unique for each giaddr assigned by a DHCP relay in the network when duplicate MAC mode is enabled.
- As is the case with unique MAC mode, client IP addresses managed by the DHCP external server application must be unique across all VRs configured on the router.
- You can configure DHCP external server to support both duplicate MAC mode (by issuing the **dhcp-external duplicate-mac-address** command) and creation of subscriber state information based on lease renewals (by issuing the **ip dhcp-external server-sync** command) simultaneously.
- DHCP external server supports the following VR topology changes for DHCP clients regardless of whether duplicate MAC mode is enabled or disabled:
 - A client roams across VRs; this might occur, for example, when a laptop computer moves to a different building in a campus network.
 - A client is assigned to a different VR; this might occur, for example, when a quality of service (QoS) policy assigns a client to a different VR during the DHCP binding process.
- When DHCP external server is configured to support unique MAC mode, which is the default, it uses only the MAC address to uniquely identify DHCP clients. Consequently, when unique MAC mode is enabled, the MAC addresses for all DHCP clients must be unique across all VRs configured on the router.

Restrictions for Using Duplicate MAC Mode to Manage Clients

The following restrictions apply when you configure the DHCP external server application to use a combination of the MAC address and giaddr to uniquely identify DHCP clients, otherwise known as duplicate MAC mode:

- You can issue the **dhcp-external duplicate-mac-address** command at any time to enable duplicate MAC mode. However, you *cannot* issue the **no dhcp-external**

duplicate-mac-address command to restore the default setting, unique MAC mode, if DHCP external server is configured for duplicate MAC mode and is currently managing any DHCP clients.

- Do not enable duplicate MAC mode for the DHCP external server application when it is configured in the same VR with either of the following:
 - An instance of the DHCP relay application that is currently managing host routes
 - Any instance of the DHCP relay proxy application
- When you enable duplicate MAC mode, the DHCP external server application ignores notifications of new clients from the RADIUS relay server application because these notifications do not include the giaddr.

**Related
Documentation**

- [Configuring DHCP External Server to Uniquely Identify Clients with Duplicate MAC Addresses on page 557](#)
- *dhcp-external duplicate-mac-address*

DHCP External Server Configuration Requirements

To configure the E Series router to support an external DHCP server, you enable the DHCP external server application on the router. If you are using DHCP packet detection, you must also specify each external DHCP server that determines which packets are monitored. The E Series router monitors all DHCP traffic between subscriber clients and the specified DHCP servers.

**Related
Documentation**

- [Configuring DHCP External Server to Uniquely Identify Clients with Duplicate MAC Addresses on page 557](#)
- [Enabling and Disabling the DHCP External Server Application on page 550](#)

Enabling and Disabling the DHCP External Server Application

Use to enable the DHCP external server application on the E Series router. Use the **no** version of the command to disable the application.

To enable the DHCP external server application on the router:

- Issue the **service dhcp-external** command:
`host1(config)#service dhcp-external`

To disable the DHCP external server application on the router:

- Issue the **no service dhcp-external** command:
`host1(config)#no service dhcp-external`

**Related
Documentation**

- *service dhcp-external*

Monitoring DHCP Traffic Between Remote Clients and DHCP Servers

You can configure the router to monitor DHCP packets between remote clients and specified DHCP servers. You can specify up to four DHCP servers.

To monitor DHCP packets between remote clients and a DHCP server:

- Issue the **ip dhcp-external server-address** command and specify the IP address of the DHCP server:

```
host1(config)#ip dhcp-external server-address 10.10.10.1
host1(config)#ip dhcp-external server-address 10.20.10.1
```

You can specify a maximum of four DHCP servers to monitor.

Related Documentation

- *ip dhcp-external server-address*

Synchronizing the DHCP External Application and the Router

In some cases the router and the DHCP external application might not be synchronized. For example, an unsynchronized condition might occur when you first enable the DHCP external server application. You can resynchronize and create subscriber state information that is based on lease renewals.

To synchronize the external DHCP server with the E Series router:

- Issue the **ip dhcp-external server-sync** command from Global Configuration mode:

```
host1(config)#ip dhcp-external server-sync
```

Related Documentation

- *ip dhcp-external server-sync*

Configuring Interoperation with Ethernet DSLAMs

The DHCP external server application uses the giaddr it receives in DHCP server-destined packets to determine the next hop for a subscriber's access routes. However, when interoperating with Ethernet digital subscriber line access multiplexers (DSLAMs), using the giaddr sent by the DSLAM can result in traffic being dropped. To ensure that traffic is forwarded properly, you can configure the DHCP external server application to disregard the DSLAM's giaddr and learn the subscriber's correct next-hop address.

The dropped traffic situation can occur because of the way some DSLAMs create the giaddr that is sent to the DHCP external server application. Some Ethernet DSLAMs use a DHCP relay implementation that inserts giaddr values and relay agent options in DHCP packets that are received from end users. The intent is that this information is provided to a DHCP server, which uses the values to determine the configuration parameters for the subscriber.

However, when the DHCP external server application receives the giaddr from an Ethernet DSLAM, the application installs the subscriber access route with the Ethernet DSLAM's IP address as the next hop. This operation results in the subscriber-destined traffic being incorrectly sent to the Ethernet DSLAM, which cannot process the traffic.

To avoid dropping the traffic in this situation, use the **ip set dhcp-external disregard-giaddr-next-hop** command to configure the DHCP external server application to ignore the giaddr when determining the next hop for the subscriber access routes. The E Series router then uses Address Resolution Protocol (ARP) to discover the subscriber's IP address and sends the traffic to the learned IP address.

To configure the DHCP external server application to ignore the giaddr when determining the next hop for the subscriber access routes:

- Issue the **ip dhcp-external disregard-giaddr-next-hop** command from Global Configuration mode:

```
host1(config)#ip dhcp-external disregard-giaddr-next-hop
```

**Related
Documentation**

- *ip dhcp-external disregard-giaddr-next-hop*

Configuring the DHCP External Server to Support the Creation of Dynamic Subscriber Interfaces

You can configure the DHCP external server to support the creation of dynamic subscriber interfaces. This configuration requires that the user's DHCP control traffic and data traffic traverse the same client-facing ingress port on the E Series router.

You must use the **ip dhcp-external auto-configure** command within a specific virtual router context.

To configure the DHCP external server to support the creation of dynamic subscriber interfaces:

- Issue the **ip dhcp-external auto-configure** command from Global Configuration mode:

```
host1(config)#ip dhcp-external auto-configure
```

To configure the DHCP external server to support the creation of dynamic subscriber interfaces built over dynamic VLANs that are based on the agent-circuit-id option (suboption 1) of the option 82 field in DHCP messages, include the **agent-circuit-identifier** keyword.

- Issue the **ip dhcp-external auto-configure** command with the **agent-circuit-identifier** keyword from Global Configuration mode:

```
host1(config)#ip dhcp-external auto-configure agent-circuit-identifier
```

The use of the option 82 field enables you to stack an IP interface that is associated with a particular subscriber over a dynamically created VLAN; the VLAN is dynamically created based on the agent-circuit-id option (suboption 1) that is contained in the DHCP option 82 field.

For information about configuring agent-circuit-id-based dynamic VLAN subinterfaces, see the *Configuring Dynamic Interfaces Using Bulk Configuration* chapter in the *JunosE Link Layer Configuration Guide*.

Related Documentation

- *ip dhcp-external auto-configure*

Configuring DHCP External Server to Control Preservation of Dynamic Subscriber Interfaces

You can configure the DHCP external server application to delete and re-create the dynamic subscriber interface after a bound client restarts the discovery process on the primary IP interface. By default, the DHCP external server preserves the existing dynamic subscriber interface in this situation.

Deleting and re-creating the dynamic subscriber interface for a DHCP client may trigger additional authentication, authorization, and accounting (AAA) services if AAA is configured on the router. Therefore, issuing the **ip dhcp-external recreate-subscriber-interface** command to delete and re-create the dynamic subscriber interface is useful if you want to use AAA services to validate and collect subscriber information during a restart of the DHCP discovery process.

Preserving the client's dynamic subscriber interface is useful if deleting and re-creating the dynamic subscriber interface might result in a service interruption.

To enable the DHCP external server application to delete and re-create the dynamic subscriber interface after a bound client restarts the discovery process on its primary IP interface:

- Issue the **ip dhcp-external recreate-subscriber-interface** command from Global Configuration mode:

```
host1:vr1(config)#ip dhcp-external recreate-subscriber-interface
```

To restore the DHCP external server default behavior to preserve the dynamic subscriber interface after a bound client restarts the discovery process on its primary IP interface:

- Issue the **no ip dhcp-external recreate-subscriber-interface** command from Global Configuration mode:

```
host1:vr1(config)#no ip dhcp-external recreate-subscriber-interface
```



NOTE: If you are upgrading the JunosE Software on the router from a release in which the DHCP external server deleted and re-created the dynamic subscriber interface by default, you must explicitly issue the **ip dhcp-external recreate-subscriber-interface** command in the current release to delete and re-create the dynamic subscriber interface. The router no longer deletes and re-creates the dynamic subscriber interface by default in this situation.

For a list of the JunosE releases in which deleting and re-creating the dynamic subscriber interface was the default behavior for the DHCP external server, see *DHCP External Server* in the *Known Behavior* section of the *JunosE Release Notes*.

**Related
Documentation**

- [Preservation of Dynamic Subscriber Interfaces with DHCP External Server Overview on page 547](#)
- *ip dhcp-external recreate-subscriber-interface*

Configuring Dynamic Subscriber Interfaces for Interoperation with DHCP Relay and DHCP Relay Proxy

When you configure the DHCP relay application or the DHCP relay proxy application in the same virtual router (VR) as the DHCP external server application, we recommend that you define interface profiles to create the dynamic subscriber interfaces when the primary IP interface is static. Otherwise, the gateway IP address (giaddr) used for DHCP rebind or renewal requests might be inconsistent. Renewal requests apply only when DHCP relay proxy is configured in the same VR as DHCP external server.

The use of inconsistent giaddrs results in the transmission of negative acknowledgment (NAK) messages from the DHCP server and the removal of existing DHCP client bindings.

To apply an interface profile to a dynamic subscriber interface:

1. Define the interface profile.

```
host1(config)#profile dsiTest
host1(config-profile)#ip unnumbered loopback 5500
host1(config-profile)#exit
```

2. Define a route map in the VR in which the static primary IP interface resides.

```
host1(config)#virtual-router relay
Proceed with new virtual-router creation? [confirm]
host1:relay(config)#route-map dsiTest
host1:relay(config-route-map)#set ip interface-profile dsiTest
host1:relay(config-route-map)#exit
```

3. Define a loopback interface with a host mask in the VR in which the static primary IP interface resides.

```
host1:relay(config)#interface loopback 5500
host1:relay(config-if)#ip address 71.23.1.2/32
```

4. Enable the static primary IP interface to support creation of dynamic subscriber interfaces, and apply the route map to the IP interface subscriber in the static primary IP interface configuration.

```
host1:relay(config)#interface fastEthernet 5/5.100
host1:relay(config-if)#ip unnumbered loopback 5500
host1:relay(config-if)#ip auto-configure ip-subscriber exclude-primary
host1:relay(config-if)#ip route-map ip-subscriber dsiTest
```

Use the **exclude-primary** keyword in the **ip auto-configure ip-subscriber** command to specify that the primary interface cannot be assigned to a subscriber.

5. If you have issued the **ip dhcp-external server-sync** command to resynchronize the DHCP external server application with the router and to support creation of subscriber state information based on lease renewals, you must do either of the following to ensure that the unicast acknowledgment (ACK) response to the renewal request has a route back to the DHCP client that generated the renewal request:

- Enable the packet detection feature on the static primary IP interface in the context of the VR in which the static primary interface resides.

```
host1:relay(config-if)#ip auto-detect ip-subscriber
```

Issuing the **ip auto-detect ip-subscriber** command creates a dynamic subscriber interface back to the DHCP client when the router receives a packet with a source IP address that does not match any entries in the demultiplexer table. This method requires you to configure the primary IP interface to support creation of dynamic subscribers interfaces, which is accomplished by issuing the **ip auto-configure ip-subscriber exclude-primary** command, as shown in Step 4.

- Configure an explicit network route in the context of the VR in which the static primary interface resides to provide connectivity back to the DHCP client.

```
host1:relay(config)#ip route 71.23.0.0/24 fastEthernet 5/5.100
```

Related Documentation

- *ip auto-configure ip-subscriber*
- *ip auto-detect ip-subscriber*
- *ip dhcp-external server-sync*

- *ip route*
- *ip route-map ip-subscriber*
- *set ip interface-profile*

Deleting Clients from a Virtual Router's DHCP Binding Table

You can delete clients from a virtual router's DHCP binding table. You can delete all clients or a specific client.



NOTE: This command is deprecated and might be removed completely in a future release. The function provided by this command has been replaced by the **dhcp delete-binding** command.

To delete clients from a virtual router's DHCP binding table, issue the **dhcp-external delete-binding** command in Privileged Exec configuration mode:

- To delete all clients:
`host1#dhcp-external delete-binding all`
- To delete a specific client:
`host1#dhcp-external delete-binding binding-id 3972819365`

Related Documentation

- *dhcp delete-binding*
- *dhcp-external delete-binding*

Configuring DHCP External Server to Uniquely Identify Clients with Duplicate MAC Addresses

You can configure the DHCP external server application to use a combination of the MAC address and giaddr to uniquely identify DHCP clients attached to the router. This behavior is also referred to as *duplicate MAC mode*. By default, DHCP external server uses only the MAC address to uniquely identify DHCP clients. The default behavior is also referred to as *unique MAC mode*.

Enabling duplicate MAC mode is useful if you are using the DHCP external server application to manage DHCP clients from multiple DHCP relays in network environments where the same MAC address might be assigned to more than one client. In such environments, DHCP external server must use a combination of the MAC address and giaddr to uniquely identify the DHCP clients it manages.

Preserving the client's dynamic subscriber interface is useful if deleting and re-creating the dynamic subscriber interface might result in a service interruption.

To configure the DHCP external server application to use a combination of the MAC address and giaddr to uniquely identify DHCP clients, also known as *enabling duplicate MAC mode*:

- Issue the **dhcp-external duplicate-mac-address** command from Global Configuration mode:

```
host1:vr1(config)#dhcp-external duplicate-mac-address
```

To restore the DHCP external server default behavior to use only the MAC address to uniquely identify DHCP clients, also known as *enabling unique MAC mode*:

- Issue the **no dhcp-external duplicate-mac-address** command from Global Configuration mode:

```
host1:vr1(config)#no dhcp-external duplicate-mac-address
```



NOTE: Unlike other commands for configuring DHCP external server, the **dhcp-external duplicate-mac-address** command applies globally to all instances of the DHCP external server application on the router, and is not issued on a per-VR basis.

Related Documentation

- [DHCP External Server Identification of Clients with Duplicate MAC Addresses Overview on page 548](#)
- `dhcp-external duplicate-mac-address`

Configuring DHCP External Server to Re-Authenticate Auto-Detected Dynamic Subscriber Interfaces

You can use the **ip re-authenticate-auto-detect ip-subscriber** command to re-authenticate the auto-detected subscribers or Dynamic Subscriber Interfaces (DSIs) created on static and dynamic primary IP interfaces, using the DHCP options when the DHCP external application manages the DSIs following a cold boot. The **no** version negates the command or restores the defaults.

To enable the IP Subscriber Manager application to re-authenticate the auto-detected subscribers created on static and dynamic primary IP interfaces after a cold boot:

- Issue the **ip re-authenticate-auto-detect ip-subscriber** command from Interface Configuration or Profile Configuration mode:

```
host1:vr1(config)#ip re-authenticate-auto-detect ip-subscriber
```

Related Documentation

- [Preservation of Dynamic Subscriber Interfaces with DHCP External Server Overview on page 547](#)
- *ip dhcp-external recreate-subscriber-interface*

CHAPTER 23

Monitoring and Troubleshooting DHCP

This chapter describes the commands you can use to monitor and troubleshoot DHCP support on E Series routers.

- [Setting Baselines for DHCP Statistics on page 560](#)
- [Monitoring Addresses Excluded from DHCP Local Server Use on page 561](#)
- [Monitoring DHCP Bindings on page 562](#)
- [Monitoring DHCP Binding Information on page 562](#)
- [Monitoring DHCP Binding Count Information on page 565](#)
- [Monitoring DHCP Binding Host Information on page 567](#)
- [Monitoring DHCP Bindings \(Displaying IP Address-to-MAC Address Bindings\) on page 570](#)
- [Monitoring DHCP Bindings \(Displaying DHCP Bindings Based on Binding ID\) on page 570](#)
- [Monitoring DHCP Bindings \(Local Server Binding Information\) on page 571](#)
- [Monitoring DHCP External Server Configuration Information on page 572](#)
- [Monitoring DHCP External Server Statistics on page 573](#)
- [Monitoring DHCP External Server Duplicate MAC Address Setting on page 574](#)
- [Monitoring DHCP Local Address Pools on page 575](#)
- [Monitoring DHCP Local Server Authentication Information on page 577](#)
- [Monitoring DHCP Local Server Configuration on page 578](#)
- [Monitoring DHCP Local Server Leases on page 579](#)
- [Monitoring DHCP Local Server Statistics on page 580](#)
- [Monitoring DHCP Option 60 Information on page 583](#)
- [Monitoring DHCP Packet Capture Settings on page 584](#)
- [Monitoring DHCP Relay Configuration Information on page 585](#)
- [Monitoring DHCP Relay Proxy Statistics on page 586](#)
- [Monitoring DHCP Relay Statistics on page 588](#)
- [Monitoring DHCP Server and DHCP Relay Agent Statistics on page 591](#)
- [Monitoring DHCP Server and Proxy Client Information on page 592](#)
- [Monitoring DHCPv6 Local Server Binding Information on page 593](#)
- [Monitoring DHCPv6 Local Server DNS Search Lists on page 594](#)

- [Monitoring DHCPv6 Local Server DNS Servers on page 594](#)
- [Monitoring DHCPv6 Local Server Prefix Lifetime on page 595](#)
- [Monitoring DHCPv6 Local Server Statistics on page 595](#)
- [Monitoring DHCPv6 Local Server Authentication Information on page 596](#)
- [Monitoring the Maximum Number of Available Leases on page 597](#)
- [Monitoring Static IP Address and MAC Address Pairs Supplied by DHCP Local Server on page 599](#)
- [Monitoring Status of DHCP Applications on page 599](#)
- [Monitoring DHCP Proxy Client Bindings on page 600](#)
- [Monitoring the Maximum Number of IPv6 Prefixes Delegated Per Interface by the DHCPv6 Local Server on page 600](#)
- [Monitoring the DHCPv6 Local Server to Send Delegated-IPv6-Prefix to SRC Client on page 602](#)

Setting Baselines for DHCP Statistics

You can use the **baseline dhcp** commands to set statistics baselines for DHCP operations. The router implements the baseline by reading and storing the statistics at the time the baseline is set and then subtracting this baseline when you retrieve baseline-relative statistics.

Use the **delta** keyword with the **show dhcp** commands to display baselined statistics.

Tasks to set a baseline for DHCP statistics are:

1. [Setting a Baseline for DHCP Relay and Relay Proxy on page 560](#)
2. [Setting a Baseline for DHCP Proxy Server Statistics on page 560](#)
3. [Setting a Baseline for DHCP External Server Statistics on page 561](#)
4. [Setting a Baseline for DHCP Local Server Statistics on page 561](#)

Setting a Baseline for DHCP Relay and Relay Proxy

To set a statistics baseline for DHCP relay and DHCP relay proxy:

- Issue the **baseline dhcp relay** command:

```
host1#baseline dhcp relay
```

There is no **no** version.

Setting a Baseline for DHCP Proxy Server Statistics

To set a baseline for DHCP proxy server statistics.

- Issue the **baseline dhcp server** command:

```
host1#baseline dhcp server
```

There is no **no** version.

Setting a Baseline for DHCP External Server Statistics

To set a baseline for DHCP external server statistics.

- Issue the **baseline ip dhcp-external** command:

```
host1#baseline ip dhcp-external
```

There is no **no** version.

Setting a Baseline for DHCP Local Server Statistics

To set a baseline for DHCP local server statistics:

- Issue the **baseline ip dhcp-local** command:

```
host1#baseline ip dhcp-local
```

There is no **no** version.

To set a baseline for DHCP local server statistics for a specific ATM, Fast Ethernet, or Gigabit Ethernet interface:

- Issue the **baseline ip dhcp-local** command with the optional **interface** keyword to specify the type of interface and interface specifier:

```
host1#baseline ip dhcp-local interface atm 3/1
```

To set a baseline for DHCPv6 local server statistics:

- Issue the **baseline ipv6 dhcpv6-local** command:

```
host1#baseline ipv6 dhcpv6-local
```

Monitoring Addresses Excluded from DHCP Local Server Use

Purpose Display addresses that have been excluded by the **ip dhcp-local excluded-address** command. The DHCP local server does not allocate excluded addresses, because they are already used by devices on the subnetwork.

Action To display excluded IP addresses:

```
host1(config)#show ip dhcp-local excluded
Dhcp Excluded Addresses
```

```
-----
      Low      High
Pool   Address Address
-----
default 10.10.1.1
default 10.10.1.5 10.10.1.30
cable2  10.10.2.1
home.com 10.10.3.1
cable4  10.10.4.1
cable5  10.10.5.1
```

Meaning [Table 112 on page 562](#) lists the **show ip dhcp-local excluded** command output fields.

Table 112: show ip dhcp-local excluded Output Fields

Field Name	Field Description
Pool	Name of the pool that contains the excluded address
Low Address	Excluded address or first address in a range of addresses
High Address	Last address in a range of addresses

Related Documentation

- [show ip dhcp-local excluded](#)

Monitoring DHCP Bindings

Tasks to monitor DHCP bindings are:

- [Monitoring DHCP Binding Information on page 562](#)
- [Monitoring DHCP Binding Count Information on page 565](#)
- [Monitoring DHCP Binding Host Information on page 567](#)
- [“Monitoring DHCP Bindings \(Displaying IP Address-to-MAC Address Bindings\)” on page 570](#)
- [“Monitoring DHCP Bindings \(Displaying DHCP Bindings Based on Binding ID\)” on page 570](#)
- [“Monitoring DHCP Bindings \(Local Server Binding Information\)” on page 571](#)

Related Documentation

- [show dhcp binding](#)
- [show dhcp count](#)
- [show dhcp host](#)
- [show ip dhcp-external binding](#)
- [show ip dhcp-external binding-id](#)
- [show ip dhcp-local binding](#)

Monitoring DHCP Binding Information

Purpose Display information for specified DHCP client bindings, with results arranged in ascending order by binding ID.



NOTE: The `show dhcp binding` command replaces the `show ip dhcp-external binding`, `show ip dhcp-external binding-id`, and `show ip dhcp-local binding` commands, which are deprecated and might be removed completely in a future release.

Action To display information about all DHCP local server bindings:

```
host1:vr1#show dhcp binding local
```

BindingId	HwAddress	Type	IpSubnet	IpAddress	State
2409734593	8000.0001.9365	local	0.0.0.0	81.3.0.2	bound
2409734595	8000.0003.9365	local	0.0.0.0	81.3.0.3	bound
2409734597	8000.0005.9365	local	0.0.0.0	81.3.0.4	bound
2409734599	8000.0007.9365	local	0.0.0.0	81.3.0.5	bound
2409734605	8000.000d.9365	local	0.0.0.0	81.3.0.8	bound
2409734607	8000.000f.9365	local	0.0.0.0	81.3.0.9	bound
2409734609	8000.0011.9365	local	0.0.0.0	81.3.0.10	bound
2409734611	8000.0013.9365	local	0.0.0.0	81.3.0.11	bound
2409734618	8000.000b.9365	local	0.0.0.0	81.3.0.7	bound
2409734619	8000.0009.9365	local	0.0.0.0	81.3.0.6	bound

The output of the **show dhcp binding** command is identical to the output of the **show dhcp host** command except for the order of the client bindings. The results of the **show dhcp binding** command are arranged in ascending order by binding ID, whereas the results of the **show dhcp host** command are arranged in ascending order by IP address.

To display information about a specific DHCP binding ID:

```
host1#show dhcp binding 3070230530
```

```
BindingId: 3070230530
HwAddress: 7000.0002.9365
IpSubnet: 0.0.0.0
IpAddress: 192.168.0.90
State: bound
Type: relay-p
Server: 192.168.15.1
Giaddr: 192.168.0.1
Lease: 3600
Remaining: 2079
IpInterface: GigabitEthernet1/0/1.101
ClientId: 45-41-48-00-01-70-00-00-02-93-65
Interface:
Relay Agent:

Agent Circuit Id: test circuit id
Agent Remote Id: test remote id
Vendor Specific: 01-02-03-04-05-06-07-08-09-0a-0b-0c-0d-0e-0f-10
Unrecognized: 11-12-13-14-15-16-17-18-19-1a-1b-1c-1d-1e-1f-20
```

To display binding information for DHCP clients with a specified interface string:

```
host1:vr2#show dhcp binding interface ip71.*4
```

BindingId	HwAddress	Type	IpSubnet	IpAddress	State
3053453315	7000.0002.9365	external	0.0.0.0	71.1.0.4	bound
3053453325	7000.000c.9365	external	0.0.0.0	71.1.0.14	bound
3053453353	7000.0016.9365	external	0.0.0.0	71.1.0.24	bound

This **show dhcp binding** command uses the * (asterisk) regular expression metacharacter in the interface string to display information for DHCP client bindings on virtual router vr2 with an IP address beginning with 71 and ending with 4. The results of the **show dhcp binding** command are arranged in ascending order by binding ID.

To display binding information for DHCP clients that match the specified circuit ID string:

```
host1:vr3#show dhcp binding circuit-id \xe3
```

BindingId	HwAddress	Type	IpSubnet	IpAddress	State
3070230529	7000.0000.9365	relay-p	0.0.0.0	71.1.0.2	bound
3070230531	7000.0002.9365	relay-p	0.0.0.0	71.1.0.4	bound
3070230535	7000.0006.9365	relay-p	0.0.0.0	71.1.0.8	bound
3070230537	7000.0008.9365	relay-p	0.0.0.0	71.1.0.10	bound
3070230539	7000.000a.9365	relay-p	0.0.0.0	71.1.0.12	bound
3070230541	7000.000c.9365	relay-p	0.0.0.0	71.1.0.14	bound
3070230543	7000.000e.9365	relay-p	0.0.0.0	71.1.0.16	bound
3070230545	7000.0010.9365	relay-p	0.0.0.0	71.1.0.18	bound
3070230547	7000.0012.9365	relay-p	0.0.0.0	71.1.0.20	bound
3070230549	7000.0014.9365	relay-p	0.0.0.0	71.1.0.22	bound
3070230553	7000.0018.9365	relay-p	0.0.0.0	71.1.0.26	bound
3070230555	7000.001a.9365	relay-p	0.0.0.0	71.1.0.28	bound
3070230557	7000.001c.9365	relay-p	0.0.0.0	71.1.0.30	bound
3070230569	7000.0016.9365	relay-p	0.0.0.0	71.1.0.24	bound
3070230572	7000.0004.9365	relay-p	0.0.0.0	71.1.0.6	bound

To specify nonprintable byte codes in the circuit ID string or remote ID string, you can use the string `\xab`, where `ab` is a hex code of the byte. This **show dhcp binding** command uses the string `\xe3` to represent byte E3 in the circuit ID string. This command displays information for the DHCP client bindings on virtual router vr3 with the specified circuit ID string, with results arranged in ascending order by binding ID.

To display information about DHCP local server bindings with a specified subnet address:

```
host1:vr1#show dhcp binding local 0.0.0.0
```

To display information about DHCP bindings with a specified IP prefix:

```
host1:vr1#show dhcp binding ip—prefix 10.1.0.0/28
```

To display information about DHCP relay proxy bindings without a lower-layer interface:

```
host1:vr1#show dhcp binding relay—proxy no-interface
```

To display binding information for DHCP clients that match the specified remote ID string:

```
host1:vr1#show dhcp binding remote-id "remote id.*even"
```

Filtering the display of DHCP client bindings by the circuit ID string or remote ID string is not supported for the DHCP external server application. DHCP external server does not store information about the agent-circuit-id suboption or agent-remote-id suboption of option 82.

Meaning Table 113 on page 564 lists the **show dhcp binding** command output fields.

Table 113: show dhcp binding Output Fields

Field Name	Field Description
BindingId	Client binding ID
HwAddress	MAC address of client

Table 113: show dhcp binding Output Fields (*continued*)

Field Name	Field Description
Type	Binding type; external (DHCP external server), local (DHCP local server), or relay-p (DHCP relay proxy)
IpSubnet	For DHCP local server bindings, the subnet of the IP address assigned to the client; 0.0.0.0 for DHCP external server and DHCP relay proxy bindings
IpAddress	IP address assigned to client
State	State of the DHCP client binding
Server	IP address of the DHCP server that allocated the client IP address
Giaddr	For DHCP relay proxy the IP address of the DHCP relay proxy; for DHCP local server bindings, the IP address of the DHCP relay that sent the packet or 0.0.0.0 if the packet comes from the client; for DHCP external server bindings, the giaddr from the DHCP packet
Lease	Total time for which the IP address is available, in seconds
Remaining	Time remaining on the current lease, in seconds
IpInterface	IP interface that is associated with the client
ClientId	DHCP Option 61 received from the client
Interface	Subinterface for DHCP local server bindings; does not apply to DHCP external server and DHCP relay proxy
Relay Agent	Indicates Relay Agent Information option (option 82)
Agent Circuit Id	Suboption 1 of the DHCP Relay Agent information option
Agent Remote Id	Suboption 2 of the DHCP relay agent information option
Vendor Specific	Suboption 9 of the DHCP relay agent information option

- Related Documentation**
- [Monitoring DHCP Binding Host Information on page 567](#)
 - *show dhcp binding*

Monitoring DHCP Binding Count Information

Purpose Display count information for DHCP client bindings and interfaces.

Action To display count information for all DHCP client bindings and interfaces:

```
host1:vr1#show dhcp count
```

Type	IpSubnet	Interfaces	Clients	Assigned Clients	Bound Clients
external	0.0.0.0	30	30	30	30
relay-p	0.0.0.0	2	30	30	30

To display count information for DHCP client bindings and interfaces with the specified interface string:

```
host1:vr2#show dhcp count interface ip71.*4
```

Type	IpSubnet	Interfaces	Clients	Assigned Clients	Bound Clients
external	0.0.0.0	3	3	3	3

This **show dhcp count** command uses the * (asterisk) regular expression metacharacter in the interface string to display information for DHCP client bindings on virtual router vr2 with an IP address beginning with 71 and ending with 4.

To display count information for DHCP client bindings and interfaces that match the specified circuit ID string:

```
host1:vr3#show dhcp count circuit-id \xe3
```

Type	IpSubnet	Interfaces	Clients	Assigned Clients	Bound Clients
relay-p	0.0.0.0	1	15	15	15

To specify nonprintable byte codes in the circuit ID string or remote ID string, you can use the string `\xab`, where *ab* is a hex code of the byte. This **show dhcp count** command uses the string `\xe3` to represent byte E3 in the circuit ID string. This command displays information for the DHCP client bindings on virtual router vr3 with the specified circuit ID string, with results arranged in ascending order by binding ID.

To display count information for DHCP local server client bindings and interfaces with a specified subnet address:

```
host1:vr1#show dhcp count local 0.0.0.0
```

To display count information for DHCP client bindings and interfaces with a specified IP prefix:

```
host1:vr1#show dhcp count ip—prefix 71.1.0.0/28
```

To display count information for DHCP relay proxy client bindings without a lower-layer interface:

```
host1:vr1#show dhcp count relay—proxy no-interface
```

To display count information for DHCP client bindings that match the specified remote ID string:

```
host1:vr1#show dhcp count remote-id "remote id.*odd"
```

Filtering the display of DHCP client bindings by the circuit ID string or remote ID string is not supported for the DHCP external server application. DHCP external server does not

store information about the agent-circuit-id suboption or agent-remote-id suboption of option 82.

Meaning Table 114 on page 567 lists the **show dhcp count** command output fields.

Table 114: show dhcp count Output Fields

Field Name	Field Description
Type	Binding type; external (DHCP external server), local (DHCP local server), or relay-p (DHCP relay proxy)
IpSubnet	For DHCP local server bindings, the subnet of the IP address assigned to the client; 0.0.0.0 for DHCP external server and DHCP relay proxy bindings
Interfaces	Number of interfaces associated with this binding type; includes the number of DHCP client bindings without a lower-layer interface, if configured
Clients	Number of DHCP clients associated with this binding type
Assigned Clients	Number of DHCP clients with an assigned IP address
Bound Clients	Number of DHCP clients in a bound state

Related Documentation

- *show dhcp count*

Monitoring DHCP Binding Host Information

Purpose Display information for specified DHCP client bindings, with results arranged in ascending order by IP address. The **show dhcp host** command displays information only for DHCP client bindings with assigned IP addresses.

Action To display information about all DHCP local server bindings:

```
host1:vr1#show dhcp host local
```

BindingId	HwAddress	Type	IpSubnet	IpAddress	State
2409734593	8000.0001.9365	local	0.0.0.0	81.3.0.2	bound
2409734595	8000.0003.9365	local	0.0.0.0	81.3.0.3	bound
2409734597	8000.0005.9365	local	0.0.0.0	81.3.0.4	bound
2409734599	8000.0007.9365	local	0.0.0.0	81.3.0.5	bound
2409734619	8000.0009.9365	local	0.0.0.0	81.3.0.6	bound
2409734618	8000.000b.9365	local	0.0.0.0	81.3.0.7	bound
2409734605	8000.000d.9365	local	0.0.0.0	81.3.0.8	bound
2409734607	8000.000f.9365	local	0.0.0.0	81.3.0.9	bound
2409734609	8000.0011.9365	local	0.0.0.0	81.3.0.10	bound
2409734611	8000.0013.9365	local	0.0.0.0	81.3.0.11	bound

The output of the **show dhcp host** command is identical to the output of the **show dhcp binding** command except for the order of the client bindings. The results of the **show**

dhcp host command are arranged in ascending order by IP address, whereas the results of the **show dhcp binding** command are arranged in ascending order by binding ID.

To display binding information for DHCP clients with a specified interface string:

```
host1:vr2#show dhcp host interface ip71.*4
```

BindingId	HwAddress	Type	IpSubnet	IpAddress	State
3053453315	7000.0002.9365	external	0.0.0.0	71.1.0.4	bound
3053453325	7000.000c.9365	external	0.0.0.0	71.1.0.14	bound
3053453353	7000.0016.9365	external	0.0.0.0	71.1.0.24	bound

This **show dhcp host** command uses the * (asterisk) regular expression metacharacter in the interface string to display information for DHCP client bindings on virtual router vr2 with an IP address beginning with 71 and ending with 4. The results of the **show dhcp host** command are arranged in ascending order by IP address.

To display binding information for DHCP clients that match the specified circuit ID string:

```
host1:vr3#show dhcp host circuit-id \xe3
```

BindingId	HwAddress	Type	IpSubnet	IpAddress	State
3070230529	7000.0000.9365	relay-p	0.0.0.0	71.1.0.2	bound
3070230531	7000.0002.9365	relay-p	0.0.0.0	71.1.0.4	bound
3070230572	7000.0004.9365	relay-p	0.0.0.0	71.1.0.6	bound
3070230535	7000.0006.9365	relay-p	0.0.0.0	71.1.0.8	bound
3070230537	7000.0008.9365	relay-p	0.0.0.0	71.1.0.10	bound
3070230539	7000.000a.9365	relay-p	0.0.0.0	71.1.0.12	bound
3070230541	7000.000c.9365	relay-p	0.0.0.0	71.1.0.14	bound
3070230543	7000.000e.9365	relay-p	0.0.0.0	71.1.0.16	bound
3070230545	7000.0010.9365	relay-p	0.0.0.0	71.1.0.18	bound
3070230547	7000.0012.9365	relay-p	0.0.0.0	71.1.0.20	bound
3070230549	7000.0014.9365	relay-p	0.0.0.0	71.1.0.22	bound
3070230569	7000.0016.9365	relay-p	0.0.0.0	71.1.0.24	bound
3070230553	7000.0018.9365	relay-p	0.0.0.0	71.1.0.26	bound
3070230555	7000.001a.9365	relay-p	0.0.0.0	71.1.0.28	bound
3070230557	7000.001c.9365	relay-p	0.0.0.0	71.1.0.30	bound

To specify nonprintable byte codes in the circuit ID string or remote ID string, you can use the string `\\xab`, where `ab` is a hex code of the byte. This **show dhcp host** command uses the string `\\xe3` to represent byte E3 in the circuit ID string. This command displays information for the DHCP client bindings on virtual router vr3 with the specified circuit ID string, with results arranged in ascending order by IP address.

To display information about DHCP external server bindings with a specified subnet address:

```
host1:vr1#show dhcp host external 0.0.0.0
```

To display information about DHCP bindings with a specified IP prefix:

```
host1:vr1#show dhcp host ip—prefix 10.2.0.0/28
```

To display information about DHCP relay proxy bindings without a lower-layer interface:

```
host1:vr1#show dhcp host relay—proxy no-interface
```

To display binding information for DHCP clients that match the specified remote ID string:

```
host1:vr1#show dhcp host remote-id "remote id.*west"
```


Filtering the display of DHCP client bindings by the circuit ID string or remote ID string is not supported for the DHCP external server application. DHCP external server does not store information about the agent-circuit-id suboption or agent-remote-id suboption of option 82.

Meaning [Table 115 on page 569](#) lists the **show dhcp host** command output fields.

Table 115: show dhcp host Output Fields

Field Name	Field Description
BindingId	Client binding ID
HwAddress	MAC address of client
Type	Binding type; external (DHCP external server), local (DHCP local server), or relay-p (DHCP relay proxy)
IpSubnet	For DHCP local server bindings, the subnet of the IP address assigned to the client; 0.0.0.0 for DHCP external server and DHCP relay proxy bindings
IpAddress	IP address assigned to client
State	State of the DHCP client binding
Server	(Detailed output only) IP address of the DHCP server that allocated the client IP address
Giaddr	(Detailed output only) For DHCP relay proxy the IP address of the DHCP relay proxy; for DHCP local server bindings, the IP address of the DHCP relay that sent the packet or 0.0.0.0 if the packet comes from the client; for DHCP external server bindings, the giaddr from the DHCP packet
Lease	(Detailed output only) Total time for which the IP address is available, in seconds
Remaining	(Detailed output only) Time remaining on the current lease, in seconds
IpInterface	(Detailed output only) IP interface that is associated with the client

Related Documentation

- [Monitoring DHCP Binding Information on page 562](#)
- *show dhcp host*

Monitoring DHCP Bindings (Displaying IP Address-to-MAC Address Bindings)

Purpose Display the mapping between the assigned IP address and the MAC address of the subscriber's computer.



NOTE: This command is deprecated and might be removed completely in a future release. The function provided by this command has been replaced by the `show dhcp binding` command.

Action To display the DHCP IP address to MAC address bindings:

```
host1#show ip dhcp-external binding
```

Dhcp External Bindings						
Hardware	Giaddr	IpAddress	Server	Lease	Expire	Interface
7000.0001.9365	91.3.0.1	91.3.0.2	10.1.2.1	3600	3175	ATM3/1.100.1

Meaning [Table 116 on page 570](#) lists the `show ip dhcp-external binding` command output fields

Table 116: show ip dhcp-external binding Output Fields

Field Name	Field Description
Hardware	MAC address of subscriber's computer
Giaddr	Gateway IP address (giaddr) in the DHCP packet received from a client
IpAddress	Subscriber client's IP address
Server	DHCP server's address
Lease	Time for which the IP address is available, in seconds
Expire	Time remaining on the current lease, in seconds
Interface	Interface that is associated with the subscriber's computer

Related Documentation

- `show ip dhcp-external binding`

Monitoring DHCP Bindings (Displaying DHCP Bindings Based on Binding ID)

Purpose Display binding information for all DHCP clients.



NOTE: This command is deprecated and might be removed completely in a future release. The function provided by this command has been replaced by the `show dhcp binding` command.

Action To display DHCP binding information:

```
host1(config)#show ip dhcp-external binding-id
```

```

      Dhcp External Binding Ids
      -----
Binding Id      Hardware      Giaddr      IpAddress
-----
3053453316     7000.0001.9365  91.3.0.1    91.3.0.2

```

Meaning Table 117 on page 571 lists the `show ip dhcp-external binding-id` command output file

Table 117: show ip dhcp-external binding-id

Field Name	Field Description
Binding Id	DHCP client binding ID option value associated with the user
Hardware	MAC address of the subscriber's computer
Giaddr	Gateway IP address (giaddr) in the DHCP packet received from a client
IpAddress	IP address assigned to the client

Related Documentation

- `show ip dhcp-external binding-id`

Monitoring DHCP Bindings (Local Server Binding Information)

Purpose Display DHCP local server binding information for DHCP local server clients. Optionally, specify an IP address or an interface to display binding information for a particular address or interface.



NOTE: This command is deprecated and might be removed completely in a future release. The function provided by this command has been replaced by the `show dhcp binding` command.

Action To display DHCP local server binding information for a specific IP address:

```
host1#show ip dhcp-local binding 192.168.1.3
```

```

Dhcp Local Bindings
-----
Address      Hardware      Lease      Interface      State
-----
192.168.1.3  11-11-22-22-33-33  (600)      fastEthernet 5/0  expired

```

To display DHCP local server binding information for a specific interface:

```
host1#show ip dhcp-local binding interface fastethernet 5/0.2
```

```

Dhcp Local Bindings
-----
Address      Hardware      Lease      Interface      State
-----
192.168.0.6  40-00-00-0b-00-01  240        fastEthernet 5/0.2  bound
192.168.0.7  40-00-00-0c-00-01  240        fastEthernet 5/0.2  bound
192.168.1.3  11-11-22-22-33-33  (600)      fastEthernet 5/0.2  expired

```

Meaning [Table 118 on page 572](#) lists the **show ip dhcp-local binding** command output fields.

Table 118: show ip dhcp-local binding Output Fields

Field Name	Field Description
Address	IP address
Hardware	MAC address of subscriber's computer
Lease	Infinite, or the number of seconds in which the IP address is available; grace period is shown in parentheses for clients in a grace period
Interface	Interface whose statistics are reported
State	Binding state; expired or released state for clients currently in the grace period

Related Documentation

- [show ip dhcp-local binding](#)

Monitoring DHCP External Server Configuration Information

Purpose Display information about the router's DHCP external server application.

Action To display DHCP external server information:

```

host1#show ip dhcp-external configuration
Dhcp External : Enabled
Auto-Configure : Enabled
Server-Sync : Enabled
Disregard-Giaddr-Next-Hop : Enabled
Detect-Agent-Circuit-Id : Disabled
Recreate-Subscriber-Interface : Enabled
Duplicate-MAC-Address : Enabled

Servers:
-----
10.1.1.1

```

10.2.1.1
10.3.1.1

Meaning Table 119 on page 573 lists the **show ip dhcp-external configuration** command output fields.

Table 119: show ip dhcp-external configuration Output Fields

Field Name	Field Description
Dhcp External	Enabled or disabled
Auto-Configure	Enabled or disabled
Server-Sync	Enabled or disabled
Disregard-Giaddr-Next-hop	Enabled or disabled
Detect-Agent-Circuit-Id	Enabled or disabled
Recreate-Subscriber-Interface	Enabled or disabled
Duplicate-MAC-Address	Enabled or disabled
Servers	DHCP servers whose traffic is monitored by the E Series router

Related Documentation • *show ip dhcp-external configuration*

Monitoring DHCP External Server Statistics

Purpose Display statistics for all external DHCP servers.

Action To display statistics for all the DHCP external servers configured on the router:

```
host1(config)#show ip dhcp-external statistics
```

```
DHCP External Statistics
Server Address 10.10.32.1
```

```
-----
      Item          Count
-----
memUsage           136
bindings            1
request             69
ack (request)      1120
renew              38611
ack (renew)        38611
nak                 42
release             68
lease expirations   0
```

Meaning Table 120 on page 574 lists the **show ip dhcp-external statistics** command output fields.

Table 120: show ip dhcp-external statistics Output Fields

Field Name	Field Description
memUsage	Memory in bytes used by DHCP server
bindings	Number of IP addresses currently assigned
request	Number of DHCP request packets
ack (request)	Number of DHCP acknowledgment packets in response to DHCP requests
renew	Number of DHCP renew packets
ack (renew)	Number of DHCP acknowledgment packets in response to DHCP renewals
nak	Number of DHCP negative acknowledgment packets
release	Number of DHCP release packets
lease expirations	Number of lease expirations

Related Documentation

- [show ip dhcp-external statistics](#)

Monitoring DHCP External Server Duplicate MAC Address Setting

Purpose Display global configuration information for the DHCP external server application. Currently, this command displays the status of the method that DHCP external server uses to uniquely identify DHCP clients with duplicate MAC addresses.

Action To display the duplicate MAC address setting for DHCP external server:

```
host1#show dhcp-external
Duplicate MAC Address: Enabled
```

Meaning [Table 121 on page 574](#) lists the **show dhcp-external** command output fields.

Table 121: show dhcp-external Output Fields

Field Name	Field Description
Duplicate MAC Address	<p>Status of the identification method for DHCP clients with duplicate MAC addresses:</p> <ul style="list-style-type: none"> • Enabled—DHCP external server uses a combination of the MAC address and giaddr to uniquely identify DHCP clients. • Disabled—(Default) DHCP external server uses only the MAC address to uniquely identify DHCP clients.

Related Documentation • *show dhcp-external*

Monitoring DHCP Local Address Pools

Purpose Display the DHCP local pool configurations.

Action To display information about the local address pool:

host1#show ip dhcp-local pool

```
*****
Pool Name - ispBoston
Pool Id - 6
Domain Name - ispBoston
Network - 10.10.0.0
Mask - 255.255.255.0
NETBIOS Node Type - 1
Lease - Days:0 Hours:0 Minutes:24 Seconds:0
Grace Period - Days:0 Hours:0 Minutes:10 Seconds:0
Grace period for released leases enabled
DNS Servers
  10.10.1.1
NETBIOS Name Servers
  10.10.1.1
  10.10.1.2
Default Routers
  10.10.1.3
Server Address - 10.10.20.8
Linked Pool - cable5
High utilization threshold - 85%
Abated utilization threshold - 75%
Current utilization - 0%
Utilization trap disabled.
Shared pool allocations - 25
```

To display information about local address pool groups:

host1#show ip dhcp-local pool groups

```
DHCP Local Server Pool Groups
There is 1 group configured
*****
Group Name: pool8_7-1-Group
  Total Addresses Available: 145
  Total Addresses In Use: 0
  High Utilization Thresh: 85%
  Abated Utilization Thresh: 75%
  Current Utilization: 0%
  Trap Enabled: no
===== Pools =====
  pool8_7-1
  pool8_7-2
  pool8_7-3
  pool8_7-4
  pool8_7-5
```

Meaning [Table 122 on page 576](#) lists the **show ip dhcp-local pool** command output fields.

Table 122: show ip dhcp-local pool Output Fields

Field Name	Field Description
Pool Name	Name of the DHCP local pool
Pool Id	ID of the pool
Domain Name	Domain name assigned to the pool
Network	Addresses that the DHCP local server can provide from the pool
Mask	Subnet mask that goes with the network address
NETBIOS Node Type	Type of NetBIOS server: 1 = Broadcast 2 = Peer-to-peer 4 = Mixed 8 = Hybrid
Lease	Time for which the supplied IP address is valid
Grace Period	Length of grace period
Grace period for released leases	Status of the grace period for released leases; enabled or disabled
DNS Servers	Address of each DNS server assigned to the pool
NETBIOS Name Servers	NetBIOS server assigned to subscribers
Default Routers	Address of default router used for subscribers
Server Address	DHCP server address that is sent to subscribers
Linked Pool	Names of any pools that are linked to this pool
High utilization threshold	Threshold at or above which the utilization trap is triggered, if the trap is enabled For linked pools, utilization is calculated based on the usage of all pools in the group
Abated utilization threshold	Threshold below which the utilization trap is triggered, if the trap is enabled For linked pools, utilization is calculated based on the usage of all pools in the group

Table 122: show ip dhcp-local pool Output Fields (*continued*)

Field Name	Field Description
Current utilization	Percentage of local address pools currently used For linked pools, utilization is calculated based on the usage of all pools in the group
Utilization trap	Status of the utilization trap which is generated when the high utilization, abated utilization, or 100% utilization is reached For linked pools, utilization is calculated based on the usage of all pools in the group
Shared pool allocations	Number of addresses allocated to shared pools
Group Name	Group name; based on the name of the original pool
Total Addresses Available	Number of addresses in the group
Total Addresses In Use	Number of addresses currently being used
Trap Enabled	Status of utilization trap, yes or no
Pools	Names of pools in the group

Related Documentation

- *show ip dhcp-local pool*

Monitoring DHCP Local Server Authentication Information

Purpose Display the DHCP local server's AAA authentication configuration information and statistics.

Action To display DHCP local server AAA authentication configuration:

```
host1#show ip dhcp-local auth config
```

DHCP Local Server Authentication Configuration

```
User-Prefix      : ERX4-Boston
Domain           : ISP1.com
Password         : to4TooL8
Virtual Router   : included
Circuit Type     : included
Circuit ID       : included
MAC Address      : excluded
Option 82        : excluded
```

DHCP Local Server DHCP Options Configuration

```
RADIUS DHCP Options : excluded
```

To display DHCP local server AAA authentication statistics:

```
host1#show ip dhcp-local auth statistics
```

DHCP Local Server Authentication Statistics

Item	Count
auth requests	10
auth request failures	0
auth grants	9
auth denies	1

Meaning [Table 123 on page 578](#) lists the **show ip dhcp-local auth** command output fields.

Table 123: show ip dhcp-local auth Output Fields

Field Name	Field Description
User-Prefix	Client's user prefix
Domain	Client's domain
Password	Password used to authenticate client
Virtual Router	Client's virtual router; excluded or included
Circuit Type	Client's circuit type; excluded or included
Circuit ID	Client's circuit ID; excluded or included
MAC Address	Client's MAC address; excluded or included
Option 82	Status of client's option 82 field; excluded or included
RADIUS DHCP Options	Status of the DHCP options returned from the RADIUS server; excluded or included
auth requests	Number of authorization requests received by this DHCP local server
auth request failures	Number of authorization requests that have failed
auth grants	Number of authorization requests that have been granted
auth denies	Number of authorization requests that have been denied

Related Documentation

- *show ip dhcp-local auth*

Monitoring DHCP Local Server Configuration

Purpose Display the DHCP local server's configuration information.

Action To display configuration settings for DHCP local server:

```
host1#show ip dhcp-local
```

```
*****
```

```
    DHCP Local Server Configuration
```

```
Mode: Standalone
```

```
SNMP Traps Enabled - no
```

```
Unique Client IDs - enabled
```

Meaning [Table 124 on page 579](#) lists the **show ip dhcp-local** command output fields.

Table 124: show ip dhcp-local Output Fields

Field Name	Field Description
Mode	DHCP local server mode, equal-access or standalone
SNMP Traps Enabled	Status of DHCP local traps support, yes or no
Unique Client IDs	Status of duplicate client ID and duplicate hardware address detection, enabled or disabled

Related Documentation

- *show ip dhcp-local*

Monitoring DHCP Local Server Leases

Purpose Display lease information for a specific IP address or for all DHCP local server leases.

Action To display information about a specific DHCP local server lease:

```
host1#show ip dhcp-local leases 192.168.0.3
```

```

                                Dhcp Local Leases
                                -----
Address      Hardware      Lease      Initiated/Renewed
-----
192.168.0.3  10-06-10-00-10-33  120       THU SEP 08 2005 08:02:11 UTC
Address      Expiration      Remaining
-----
192.168.0.3  THU SEP 08 2005 08:04:11 UTC  79
Address      Initial Lease Start
-----
192.168.0.3  THU SEP 08 2005 08:01:12 UTC

```

To display information about all DHCP local server leases:

```
host1#show ip dhcp-local leases
```

```

                                Dhcp Local Leases
                                -----
Address      Hardware      Lease      Initiated/Renewed
-----
192.168.0.2  10-06-10-00-10-32  120       THU JUL 06 2006 08:02:11 UTC
192.168.0.3  10-06-10-00-10-33  120       THU JUL 06 2006 08:02:11 UTC
192.168.55.4 10-06-10-00-10-34  (600)     THU JUL 06 2006 09:57:22 UTC

```

```

192.168.55.5  10-06-10-00-10-35  infinite  THU JUL 06 2006 08:03:10 UTC
Address      Expiration      Remaining
-----
192.168.0.2   THU JUL 06 2006 08:04:11 UTC  80
192.168.0.3   THU JUL 06 2006 08:04:11 UTC  80
192.168.55.4   THU JUL 06 2006 10:07:22 UTC  575
192.168.55.5   THU JUL 06 2006 08:04:11 UTC  infinite
Address      Initial Lease Start
-----
10.1.0.2      THU JUL 06 2006 08:01:12 UTC
10.1.0.3      THU JUL 06 2006 08:01:12 UTC
192.168.55.4   THU JUL 06 2006 09:54:19 UTC
192.168.55.5   THU JUL 06 2006 08:03:10 UTC

```

Meaning [Table 125 on page 580](#) lists the **show ip dhcp-local leases** command output fields.

Table 125: show ip dhcp-local leases Output Fields

Field Name	Field Description
Address	IP address
Hardware	MAC address of the subscriber's computer
Lease	Infinite, or the number of seconds in which the IP address is available; grace period in parentheses for clients in the grace period
Initiated/Renewed	Day, date, and time the lease was most recently initiated or renewed; start time of grace period for clients in the grace period
Expiration	Day, date, and time the lease expires; expiration time of grace period for clients in the grace period
Remaining	Infinite, or the number of seconds remaining in the lease, if any; remaining time of grace period for clients in the grace period
Initial Lease Start	Day, date, and time the lease was initiated

Related Documentation

- *show ip dhcp-local leases*

Monitoring DHCP Local Server Statistics

Purpose Display statistics for the DHCP local server.

Action To display all DHCP local server statistics:

```
host1#show ip dhcp-local statistics
```

```
DHCP Local Server Statistics
```

```

-----
Item      Count
-----
memUsage  184
bindings  2

```

```

--Receive Statistics--
discover                8
request(accept)         10
request(renew)           6
request(rebind)          2
request(other)           6
decline                  0
release                  6
inform                   0
total in packet          38
in error                  0
in discard                0
unknown client packet    6
--Transmit Statistics--
offer                    8
ack(accept)              10
ack(renew)                6
ack(rebind)               2
nak                       6
nak(renew)                0
nak(rebind)               0
total out packet          32
out error                 0
out discard                0

```

To display DHCP local server statistics for a specific interface:

```
host1#show ip dhcp-local statistics interface atm 4/0.32
```

DHCP Local Server SubInterface Statistics		
Interface	Item	Count

ATM4/0.32		
Receive Statistics		
	discover	4
	request(accept)	5
	request(renew)	1
	request(rebind)	1
	request(other)	3
	decline	0
	release	3
	inform	0
	total in packet	17
	in error	0
	in discard	0
	unknown client packet	3
Transmit Statistics		
	offer	4
	ack(accept)	5
	ack(renew)	1
	ack(rebind)	1
	nak	3
	nak(renew)	0
	nak(rebind)	0
	total out packet	14
	out error	0
	out discard	0

Meaning [Table 126 on page 582](#) lists the `show ip dhcp-local statistics` command

Table 126: show ip dhcp-local statistics output fields

Field Name	Field Description
memUsage	Number of bytes of memory used by the DHCP local server
bindings	Number of leased IP addresses currently assigned
Receive Statistics	Statistics for packets that have been received
discover	Number of DHCP discover messages received
request(accept)	Number of DHCP requests accepted
request(renew)	Number of DHCP requests for renewal received
request(rebind)	Number of DHCP requests for rebinding received
request(other)	Number of DHCP unknown requests received
decline	Number of DHCP decline messages received
release	Number of DHCP release messages received
inform	Number of DHCP inform messages received
total in packet	Number of packets received
in error	Number of packets received with errors that prevent further processing; count is independent of the message-type counters
in discard	Number of packets received that are discarded due to system resource issues; count is independent of the message-type counters
unknown client packet	Number of nonrequest packets that have no entry in the local server database received
Transmit Statistics	Statistics for packets that have been transmitted
offer	Number of DHCP offer messages sent
ack(accept)	Number of DHCP acknowledgments sent in response to accepted requests
ack(renew)	Number of DHCP acknowledgments sent in response to renewal requests
ack(rebind)	Number of DHCP acknowledgments sent in response to rebinding requests
nak	Number of DHCP NAK messages sent in response to requests that cannot be bound or that are unknown to this local server

Table 126: show ip dhcp-local statistics output fields (*continued*)

Field Name	Field Description
nak(renew)	Number of DHCP NAK messages sent in response renewal requests
nak(rebind)	Number of DHCP NAK messages sent in response to rebinding requests
total out packet	Number of packets sent by the DHCP local server
out error	Number of packets that cannot be transmitted due to protocol errors or configuration errors; count is independent of the message-type counters
out discard	Number of packets that cannot be transmitted due to system resource issues; count is independent of the message-type counters

Related Documentation

- *show ip dhcp-local statistics*

Monitoring DHCP Option 60 Information

- Purpose** Display configuration and action information for the DHCP vendor-option (option 60) feature.
- Use the command without additional keywords to display information for all vendor option configurations.
 - Use the **vendor-option-relay-server** keyword and server address to display information for option 60 strings that match a configured string that results in the packets being sent to the specified vendor-option server.
 - Use the **default** keyword to display information for option 60 strings that do not match a configured vendor-option string.

Action To display information for all vendor option configurations:

host1#show dhcp vendor-option

Codes:

- * - the configured vendor-string is an exact-match
- default - all DHCP client packets not matching a configured vendor-string
- implied - the DHCP application is configured but has not been enabled with the vendor-option command
- drop - the DHCP application responsible for the action has not been configured yet therefore all packets for this application will be dropped

Total 4 entries.

Vendor-option	Action
Juniper	relay to 10.10.1.1 (rx: 0)
default(*)	relay to 192.168.5.5 (rx: 0, no-match: 0)
someString(*)	relay to 192.168.7.7 (rx: 0)
someString2(*)	local-server (rx: 0)

To display information for option 60 strings that match a configured string:

```
host1#show dhcp vendor-option vendor-option-relay-server 10.10.1.1
```

Codes:

- * - the configured vendor-string is an exact-match
- default - all DHCP client packets not matching a configured vendor-string
- implied - the DHCP application is configured but has not been enabled with the vendor-option command
- drop - the DHCP application responsible for the action has not been configured yet therefore all packets for this application will be dropped

Total 4 entries.

Vendor-option	Action
Juniper	relay to 10.10.1.1 (rx: 0)

Meaning [Table 127 on page 584](#) lists the **show dhcp vendor-option** command output fields.

Table 127: show dhcp vendor-option Output Fields

Field Name	Field Description
Vendor-option	Option 60 string; an asterisk (*) indicates that the string exactly matches a configured option 60 string, default indicates the action to take when the string does not match a configured option 60 string
Action	Action to take for the indicated string match; drop, forward to local-server, proxy client server, or all configured DHCP vendor option servers; or relay to the specified DHCP server
rx	Received packets that match a vendor-option string
no-match	Received packets that do not match a vendor-option string; no-match statistics appear only for default entries

Related Documentation

- *show dhcp vendor-option*

Monitoring DHCP Packet Capture Settings

Purpose Display the configuration for per-interface DHCP packet logging.

Action To display configuration information about the DHCP packet capture feature:

```
host1#show ip dhcp-capture
```

```

      Dhcp Capture Configuration
      -----
Router  Interface  Type   Priority
-----
default ip3/1    Rx/Tx  low/low
default ip5/1    Rx     high

```

Meaning [Table 128 on page 585](#) lists the **show ip dhcp-capture** command output fields.

Table 128: show ip dhcp-capture Output Fields

Field Name	Field Description
Router	Router name
Interface	Interface whose DHCP packets are logged
Type	Packet type to be logged, Rx (received), Tx (transmitted), or Rx/Tx (all)
Priority	Priority assigned to logged packets, low or high

Related Documentation

- *show ip dhcp-capture*

Monitoring DHCP Relay Configuration Information

Purpose Display DHCP relay configuration information and the IP addresses of the configured DHCP servers.

Action To display information about the DHCP relay configuration and the IP address of the DHCP servers.

```
host1#show dhcp relay
DHCP Relay Configuration
-----
Mode: Proxy
  Restore Client Timeout: 72
  Send First Offer: off
Inhibit Access Route Creation: off
Assign Giaddr to Source IP: off
Layer 2 Unicast Replies: off
Giaddr Selects Interface: off
Broadcast Flag Replies: on
Maximum client pps : 4096
Relay Agent Information Option (82):
  Override Giaddr: off
  Override Option: off
  Trust All Clients: off
  Preserve Option From Trusted Clients: off
Circuit-ID Sub-option (1): on
  select - hostname
  select - exclude-subinterface-id
Remote-ID Sub-option (2): on
Vendor-Specific Sub-option (9): on
  select - layer2-circuit-id
  select - user-packet-class

DHCP Server Addresses
-----
30.3.7.1
```

Meaning Table 129 on page 586 lists the **show dhcp relay** command output fields.

Table 129: show dhcp relay Output Fields

Field Name	Field Description
Mode	DHCP relay mode; either Standard (DHCP relay mode) or Proxy (DHCP relay proxy mode)
Restore Client Timeout	(DHCP relay proxy mode only) number of hours
Send First Offer	On or off
Inhibit Access Route Creation	On or off
Assign Giaddr to Source IP	On or off
Layer 2 Unicast Replies	On or off
Giaddr Selects Interface	On or off
Broadcast Flag Replies	On or off
Maximum client pps	Maximum number of client packets processed per second
Override Giaddr	On or off
Override Option	On or off
Trust All Clients	On or off
Preserve Option From Trusted Clients	On or off
Circuit-ID Sub-option (1)	On or off; when on includes a list of selected suboptions
Remote-ID Sub-option (2)	On or off
Vendor-Specific Sub-option (9)	On or off; when on includes a list of selected suboptions
DHCP Server Addresses	IP addresses of configured DHCP servers

Related Documentation

- *show dhcp relay*

Monitoring DHCP Relay Proxy Statistics

Purpose Display statistics for the DHCP relay proxy.



NOTE: The `show dhcp relay statistics` command displays additional DHCP statistics that the router reports for both DHCP relay and DHCP relay proxy.

Action To display DHCP relay proxy statistics:

```
host1# show dhcp relay proxy statistics
```

DHCP Relay/Proxy Statistics

Address	Disc.	Offer	Req.	Ack	Nak	Decline	Release	Inform
192.168.1.1	9	0	0	0	0	0	0	0
192.168.1.2	9	0	0	0	0	0	0	0
192.168.32.1	9	5	5	5	0	0	0	0

Active Clients: 5

Clients to Restore: 0

Client Packets: 14

Server Packets: 10

Timed Out: 0

No Offers: 4

Modify Fail: 0

Meaning [Table 130 on page 587](#) lists the `show dhcp relay proxy statistics` command output fields.

Table 130: show dhcp relay proxy statistics Output Fields

Field Name	Field Description
Address	IP address of the DHCP server
Disc.	Number of discover messages sent to server
Offer	Number of offers received from a server
Req.	Number of requests sent to a server
Ack	Number of ACK messages received from a server
Nak	Number of NAK messages received from a server
Decline	Number of decline messages sent to a server
Release	Number of releases sent to a server
Inform	Number of information messages sent to a server
Active Clients	Number of clients being maintained by the relay proxy
Clients to Restore	Number of host routes installed without an active client (waiting for renewal)
Client Packets	Total number of packets received from clients

Table 130: show dhcp relay proxy statistics Output Fields (*continued*)

Field Name	Field Description
Server Packets	Total number of packets received from servers
Timed Out	Number of clients removed because of lease expiration
No Offers	Number of clients removed because no server sent an offer
Modify Fail	Number of clients deleted because the relay proxy failed to modify the DHCP packet

Related Documentation

- *show dhcp relay proxy statistics*

Monitoring DHCP Relay Statistics

Purpose Display DHCP packet error and relay agent option statistics that are reported for both DHCP relay and DHCP relay proxy, and also to display DHCP server statistics related only to DHCP relay.



NOTE: The `show dhcp relay proxy statistics` command displays additional DHCP statistics that the router reports only for DHCP relay proxy.

Action To display DHCP relay statistics:

```
host1# show dhcp relay statistics
```

```

DHCP Relay Statistics
-----
Statistic                                     Values
-----
Packet error statistics (standard & proxy modes):
  dropped discover packets, no resources      0
  dropped dhcp packets, no resources          0
  dropped bad message operation packets       0
  dropped unknown message type request packets 0
  dropped unknown message type reply packets  0
Packet Pacing Algorithm (standard & proxy modes):
  Speed up pacer                             0
  Slow down pacer                           0
  Current client pps                         1024
Relay Agent Option statistics (standard & proxy modes):
  add Relay Agent Option circuit ID suboption On
  add Relay Agent Option remote ID suboption On
  packets with giaddr override               0
  packets with Relay Agent Option override   2
  packets forwarded with Relay Agent Option already present 4
  dropped packets with Relay Agent Option already present 3
  dropped giaddr spoof packets               0
DHCP server statistics (standard mode only):
  dropped duplicate request packets          12

```

```

packets transmitted to servers          38
packets received from servers          26
dropped unknown xid reply packets      0
dropped stale request packets          12

```

To display detailed statistics for DHCP relay only, use the optional **detail** keyword—this command displays DHCP server statistics and dropped unknown message type reply packets statistics:

```
host1# show dhcp relay statistics detail
```

```

                                DHCP Relay Detail Statistics
                                -----
                                Statistics      10.10.1.1    192.168.32.12    192.168.32.1
                                -----
Dropped unknown message type replies  0             0             0
Dropped duplicate requests             6             6             0
Packets transmitted to server          6             6            26
Packets received from server           0             0            26
Dropped unknown xids replies           0             0             0
Dropped stale requests                 6             6             0

```

Meaning Table 131 on page 589 lists the **show dhcp relay statistics** command output fields.

Table 131: show dhcp relay statistics Output Fields

Field Name	Field Description
Packet error statistics (standard & proxy modes)	
dropped discover packets, no resources	Number of received DHCP relay discover messages that were discarded because of lack of resources
dropped dhcp packets, no resources	Number of received DHCP relay messages, other than discover messages that were discarded because of lack of resources
dropped bad message operation packets	Number of received DHCP relay messages that were discarded because their message operation (for example, bootrequest, bootreply) was unknown, possibly due to corruption
dropped unknown message type request packets	Number of received DHCP relay request messages that were discarded because their message type (for example, discover, offer-request) was unknown, possibly due to corruption
dropped unknown message type reply packets	Number of received DHCP relay reply messages that were discarded because their message type (for example, offer, ack) was unknown, possibly due to corruption
Packet Pacing Algorithm (standard & proxy modes)	
Speed up pacer	Number of times the DHCP relay increased the rate of client packets processed

Table 131: show dhcp relay statistics Output Fields (*continued*)

Field Name	Field Description
Slow down pacer	Number of times the DHCP relay decreased the rate of client packets processed
Current client pps	Rate of client packets processed per second
Relay Agent Option statistics (standard & proxy modes)	
add Relay Agent Option circuit ID suboption	Status of circuit ID suboption, on or off
add Relay Agent Option remote ID suboption	Status of remote ID suboption, on or off
packets with giaddr override	Number of received DHCP relay requests whose giaddr field is overridden with IP address 0.0.0.0
packets with Relay Agent Option override	Number of received DHCP relay requests whose relay agent information option is overridden with an option string created by this relay agent
packets forwarded with Relay Agent Option already present	Number of received DHCP relay requests already containing the relay agent information option that were forwarded to DHCP servers
dropped packets with Relay Agent Option already present	Number of received DHCP relay requests that were discarded because they already contained the relay agent information option when this relay agent was configured to insert the option
dropped giaddr spoof packets	Number of received DHCP relay requests that were discarded because the gateway IP address field already contained this relay agent's IP address
DHCP server statistics (standard mode only)	
dropped duplicate request packets	Number of received DHCP relay requests that were discarded because they have a matching server address and XID of an outstanding DHCP server request
packets transmitted to servers	Number of DHCP relay requests successfully transmitted to DHCP servers
packets received from servers	Number of DHCP relay replies successfully received from DHCP servers
dropped unknown xid reply packets	Number of DHCP relay replies received from DHCP servers that were discarded because their server address and XID do not match an outstanding DHCP server request

Table 131: show dhcp relay statistics Output Fields (*continued*)

Field Name	Field Description
dropped stale request packets	Number of DHCP relay requests sent to DHCP servers that were discarded because their replies timed out

Related Documentation

- *show dhcp relay statistics*

Monitoring DHCP Server and DHCP Relay Agent Statistics

Purpose Display DHCP proxy server statistics

Action To display statistics for the DHCP proxy server:

```
host1#show dhcp server statistics
DHCP Proxy Global Statistics
  Messages from Unknown Servers 0

DHCP Proxy Server Statistics
-----
```

Statistic	Counts	Counts	Counts
DHCP Server Address	10.6.128.10	10.10.0.42	192.168.200.10
Discovers sent	0	0	0
Leases granted	0	0	0
Offers received	0	0	0
Requests sent	0	0	0
Acks received	0	0	0
Naks received	0	0	0
addresses declined	0	0	0
addresses released	0	0	0
Inform sent	0	0	0
unknown messages	0	0	0
bad messages	0	0	0

Meaning [Table 132 on page 591](#) lists the **show dhcp server statistics** command output fields

Table 132: show dhcp server statistics Output Fields

Field Name	Field Description
DHCP Server Address	IP address of the server
Discovers sent	Number of discover messages sent by the server
leases granted	Number of leases granted by the server
Offers received	Number of offers sent by the server
Requests sent	Number of requests sent to the server
Acks received	Number of acknowledgments received from the server

Table 132: show dhcp server statistics Output Fields (*continued*)

Field Name	Field Description
Naks received	Number of negative acknowledgments received from the server
addresses declined	Number of IP addresses rejected because they were already in use
addresses released	Number of IP addresses released back to the server
Inform sent	Number of inform messages sent to the server
unknown messages	Number of illegal DHCP messages or messages that cannot be handled by the router
bad messages	Number of messages not recognized as DHCP messages

Related Documentation

- *show dhcp server statistics*

Monitoring DHCP Server and Proxy Client Information

Purpose Display DHCP server and proxy client information.

Action To display information about the DHCP server and proxy clients:

```
host1#show dhcp server
```

```

DHCP Proxy Client Status:
-----
O A Address      Leases Offers Requests Acks Naks Declines Releases
- - - - -
E E 10.6.128.10   0      0      0      0      0      0      0
E E 10.6.128.11   0      0      0      0      0      0      0

```

Meaning [Table 133 on page 592](#) lists the **show dhcp server** command output fields.

Table 133: show dhcp server Output Fields

Field Name	Field Description
O	Read-only value that displays the operational status of the server
A	Read/write value that displays the administrative status of the server
E	Enabled; indicates that the server is being actively used to supply IP addresses to clients
D	Draining; indicates that the server is not accepting any new requests for addresses, but is maintaining the addresses that it has already assigned
X	Disabled; means that the server is not accepting any new requests for addresses and has no outstanding addresses

Table 133: show dhcp server Output Fields (*continued*)

Field Name	Field Description
Address	IP address of a DHCP server
Leases	Number of IP address leases granted by the server
Offers	Number of offers sent by the server
Requests	Number of requests sent to the server
Acks	Number of acknowledgments received from the server
Naks	Number of negative acknowledgments received from the server
Declines	Number of IP addresses rejected because they were already in use
Releases	Number of IP addresses released back to the server

Related Documentation

- *show dhcp server*

Monitoring DHCPv6 Local Server Binding Information

Purpose Display the mapping between one or more IPv6 addresses and the DHCP unique ID of the subscriber's computer.

Action To display the DHCP binding information for an IPv6 address:

```
host1#show ipv6 dhcpv6-local binding 2001:db8:4::/48
```

Prefix	Client DUID	Lease	Intf
-----	-----	-----	-----
2001:db8:4::/48	<LL 1/00A0DE113502>	infinite	FastEthernet 3/6.1

Meaning [Table 134 on page 593](#) lists the **show ipv6 dhcpv6-local binding** command output fields.

Table 134: show ipv6 dhcpv6-local binding Output Fields

Field Name	Field Description
Prefix	IPv6 address
Client DUID	DHCP unique ID of subscriber's computer
Lease	Time for which the IPv6 address is available in seconds, or infinite
Intf	Router's interface that is associated with the subscriber's computer

Related Documentation

- *show ipv6 dhcpv6-local binding*

Monitoring DHCPv6 Local Server DNS Search Lists

Purpose Display the DHCPv6 local servers DNS search list.

Action To display the DNS search list for DHCPv6 local servers:

```
host1#show ipv6 dhcpv6-local dns-domain-searchlist
Domain 1: xyzcorporation.net
Domain 2: xyzcorp.com
Domain 3: financeDomain.com
Domain 4: researchDomain.com
```

Meaning [Table 135 on page 594](#) lists the **show ipv6 dhcpv6-local dns-domain-searchlist** command output fields.

Table 135: show ipv6 dhcpv6-local dns-domain-searchlist Output Fields

Field Name	Field Description
Domain	Domains in the search list

Related Documentation

- *show ipv6 dhcpv6-local dns-domain-searchlist*

Monitoring DHCPv6 Local Server DNS Servers

Purpose Display a list of DNS servers configured on the DHCPv6 local server.

Action To display the list of DNS servers:

```
host1#show ipv6 dhcpv6-local dns-servers
DNS server 1: 2001:db8:18::
DNS server 2: 2001:db8:19::
DNS server 3: 2001:db8:20::
DNS server 4: 2001:db8:21::
```

Meaning [Table 136 on page 594](#) lists the **show ipv6 dhcpv6-local dns-servers** command output fields.

Table 136: show ipv6 dhcpv6-local dns-servers Output Fields

Field Name	Field Description
DNS server	IPv6 address of the DNS server

Related Documentation

- *show ipv6 dhcpv6-local dns-servers*

Monitoring DHCPv6 Local Server Prefix Lifetime

Purpose Display the DHCPv6 default prefix lifetime.

Action To display the DHCPv6 default prefix lifetime:

```
host1#show ipv6 dhcpv6-local prefix-lifetime
default prefix lifetime is 1 day, 12 hours, 30 minutes
```

Meaning [Table 137 on page 595](#) lists the `show ipv6 dhcpv6-local prefix-lifetime` command output fields.

Table 137: show ipv6 dhcpv6-local prefix-lifetime Output Fields

Field Name	Field Description
default prefix lifetime	Number of days, hours, and minutes

Related Documentation

- *show ipv6 dhcpv6-local prefix-lifetime*

Monitoring DHCPv6 Local Server Statistics

Purpose Display statistics for the DHCPv6 local server.

Action To display DHCPv6 local server statistics:

```
host1#show ipv6 dhcpv6-local statistics
```

```
DHCPv6 Local Server Statist
```

```
-----
      Item              Count
-----
memUsage                136
bindings                 1
solicit rx              1
request(accept) rx      1
request(renew) rx       0
decline rx              0
release rx              0
inform rx               0
confirm rx              0
rebind rx               0
reconfigure tx          0
advertise tx            1
successful reply tx     1
failed reply tx         0
unknown msgs            0
bad msgs                0
```

Meaning [Table 138 on page 596](#) lists the `show ipv6 dhcpv6-local statistics` command output fields.

Table 138: show ipv6 dhcpv6-local statistics Output Fields

Field Name	Field Description
memUsage	Number of bytes of memory used by DHCPv6 local server
bindings	Number of leased IPv6 prefixes currently assigned
solicit rx	Number of DHCPv6 solicit messages received
request(accept) rx	Number of DHCPv6 request messages received
request(renew) rx	Number of DHCPv6 requests for renewal received
decline rx	Number of DHCPv6 decline messages received
release rx	Number of DHCPv6 release messages received
inform rx	Number of DHCPv6 information-request messages received
confirm rx	Number of DHCPv6 confirm messages received
rebind rx	Number of DHCPv6 rebind messages received
reconfigure tx	Number of DHCPv6 reconfigure messages transmitted
advertise tx	Number of DHCPv6 advertise messages transmitted
successful reply tx	Number of reply messages transmitted with success reply code
failed reply tx	Number of reply messages transmitted with reply codes other than success
unknown msgs	Unused field; always 0
bad msgs	Number of messages with errors received by the DHCPv6 local server

Related Documentation

- *show ipv6 dhcpv6-local statistics*

Monitoring DHCPv6 Local Server Authentication Information

Purpose Display the DHCPv6 local server's AAA authentication configuration information.

Action To display the DHCPv6 local server's AAA authentication configuration:

```
host1#show ipv6 dhcpv6-local auth config
```

DHCPv6 Local Server Authentication Configuration

User-Prefix : userPrefix

Domain : domain

```

Password          : password
Circuit Type      : excluded
Circuit ID        : excluded

```

Meaning [Table 139 on page 597](#) lists the **show ipv6 dhcpv6-local auth config** command output fields.

Table 139: show ipv6 dhcpv6-local auth config Output Fields

Field Name	Field Description
User-Prefix	Client's user prefix
Domain	Client's domain
Password	Password used to authenticate client
Circuit Type	Client's circuit type; excluded or included
Circuit ID	Client's circuit ID; excluded or included

Related Documentation

- [Authentication and Accounting of IPv6 Subscribers Using the DHCPv6 Local Server Overview on page 480](#)
- [Interoperation of Authentication of IPv6 Clients and Display of Active Subscriber Information on page 482](#)
- [Configuring AAA Authentication for DHCPv6 Local Server Standalone Mode on page 500](#)
- *show ipv6 dhcpv6-local auth config*

Monitoring the Maximum Number of Available Leases

Purpose Display the maximum number of leases available for each VPI/VCI, VLAN, Ethernet subnetwork, or POS access interface type, or for a specific interface or subinterface.

Action To display the maximum number of leases available for each interface type:

```
host1(config)#show ip dhcp-local limits
```

```

*****
DHCP Local Server Address Limits
ATM Limit      - 48000
VLAN Limit     - 48000
POS Limit      - 1000
Ethernet Limit - 48000

```

To display information about the maximum number of leases for a specific interface:

```
host1(config)#show ip dhcp-local limits interface atm 3/1
```

```

DHCP Local Interface Limits
-----
Interface  Limit  Count  Denied  Total
                        Denied

```

```

-----
atm 3/1      300      127      5      29

```

To display information about the maximum number of leases on all interfaces:

```
host1(config)#show ip dhcp-local limits interface
```

```

          Dhcp Local Interface Limits
          -----
Interface      Limit    Count    Denied    Total
-----
fastEthernet0/0 200      0        0         0
atm 3/1        300     127      5         29
atm 4/2        5000     0        0         0
atm 5/1        5000    15        2         5
pos 2/1        1000     0        0         0

```

Meaning [Table 140 on page 598](#) lists the **show ip dhcp-local limits** command output fields.

Table 140: show ip dhcp-local limits Output Fields

Field Name	Field Description
ATM Limit	Number of leases available for each VPI/VCI
VLAN Limit	Number of leases available for each VLAN
POS Limit	Number of leases available for each POS access interface
Ethernet Limit	Number of leases available for each Ethernet subnet
Limit	Number of leases available to the specified interface or subinterface; indicates the configured value for the interface type unless a specific lease value is configured for the particular interface
Count	Number of active leases on the interface
Denied	Number of lease requests denied during the current denial period; this number is reset to zero (and the denial period restarted) when the number of active leases no longer exceeds the configured limit
Total Denied	Total number of lease requests denied on the interface since the interface became active

Related Documentation

- show ip dhcp-local limits*

Monitoring Static IP Address and MAC Address Pairs Supplied by DHCP Local Server

Purpose Display the static IP address/MAC address pairs that the DHCP local server supplies in standalone mode.

Action To display information about static IP address/MAC address pairs:

```
host1#show ip dhcp-local reserved
```

```

                                Dhcp Reserved Addresses
                                -----
                                Pool          Address          Hardware
                                -----
cablemodem 10.44.44.100 12-34-12-34-12-34-00-00-00-00-00-00-00-00-00-00
cablemodem 10.44.44.101 22-33-22-33-22-33-00-00-00-00-00-00-00-00-00-00

```

Meaning [Table 141 on page 599](#) lists the **show ip dhcp-local reserved** command output fields.

Table 141: show ip dhcp-local reserved Output Fields

Field Name	Field Description
Pool	Name of pool in which the address is reserved
Address	IP address that is reserved
Hardware	Address for which the IP address is reserved

Related Documentation

- *show ip dhcp-local reserved*

Monitoring Status of DHCP Applications

Purpose Display which DHCP applications are configured whether they are active or inactive—displays the status of DHCP relay, DHCP relay proxy, DHCP local server, and DHCP external server.

Action To display the status of the configured DHCP applications:

```

host1#show dhcp summary
DHCP local-server configured and inactive
DHCP relay configured and active

```

Meaning [Table 142 on page 599](#) lists the **show dhcp summary** command output fields.

Table 142: show dhcp summary Output Fields

Field Name	Field Description
configured	Applications that are currently configured
active or inactive	Current status of the application

Related Documentation • [show dhcp summary](#)

Monitoring DHCP Proxy Client Bindings

Purpose Display information for all DHCP proxy client bindings, with results arranged in ascending order by binding ID.

Action To display information for all DHCP proxy client bindings:

```
host1#show dhcp proxy-client binding
Dhcp Proxy Client Bindings
-----
```

ClientId	IpAddress	Server	Lease	Expire	State
436600832	21.1.0.1	31.1.1.2	600	466	bound
436600833	21.1.0.2	31.1.1.2	600	467	bound
436600834	21.1.0.3	31.1.1.2	600	468	bound
436600835	21.1.0.4	31.1.1.2	600	468	bound
436600836	21.1.0.5	31.1.1.2	600	468	bound
436600837	21.1.0.6	31.1.1.2	600	468	bound
436600838	21.1.0.7	31.1.1.2	600	468	bound
436600839	21.1.0.8	31.1.1.2	600	468	bound
436600840	21.1.0.9	31.1.1.2	600	468	bound
436600841	21.1.0.10	31.1.1.2	600	468	bound

Meaning [Table 143 on page 600](#) lists the **show dhcp proxy-client binding** command output fields.

Table 143: show dhcp proxy-client binding Output Fields

Field Name	Field Description
ClientId	Proxy client binding ID
IpAddress	IP address assigned to the client
Server	IP address of the DHCP server that allocated the client IP address
Lease	Total time for which the IP address is available, in seconds
Expire	Time remaining on the current lease, in seconds
State	State of the DHCP client binding

Related Documentation • [show dhcp proxy-client binding](#)

Monitoring the Maximum Number of IPv6 Prefixes Delegated Per Interface by the DHCPv6 Local Server

Purpose Determine the maximum number of IPv6 prefixes that can be delegated by the DHCPv6 local server for ATM, Ethernet, VLAN, or POS interface types, or for a particular interface or subinterface.

Action To determine the maximum number of IPv6 prefixes that can be delegated for ATM, Ethernet, VLAN, or POS interface types, enter the following command:

```
host1#show ipv6 dhcpv6-local limits
```

```
*****
```

DHCPv6 Local Server Prefix Limits

```
ATM Limit      - 48000
VLAN Limit     - 48000
POS Limit      - 48000
Ethernet Limit - 48000
```

To determine the maximum number of IPv6 prefixes that can be delegated by the DHCPv6 local server to DHCPv6 subscribers per interface and to view the subinterface table details for a particular interface or subinterface, enter the following command:

```
host1#show ipv6 dhcpv6-local limits interface gigabitEthernet 6/0
```

Dhcpv6 Local Interface Limits

```
-----
Interface      Limit  Count  Denied  Total
-----
GigabitEthernet6/0  500    0      0      0
```

To determine the maximum number of IPv6 prefixes that can be delegated by the DHCPv6 local server to DHCPv6 subscribers per interface and to view the subinterface table details for all interfaces, enter the following command:

```
host1#show ipv6 dhcpv6-local limits interface
```

Dhcpv6 Local Interface Limits

```
-----
Interface      Limit  Count  Denied  Total
-----
GigabitEthernet6/0  500    0      0      0
ATM5/0.100       1000   0      0      0
```

Meaning [Table 144 on page 601](#) lists the `show ipv6 dhcpv6-local limits` command output fields.

Table 144: show ipv6 dhcpv6-local limits Output Fields

Field Name	Field Description
ATM Limit	Number of IPv6 prefixes that can be delegated for each VPI or VCI
VLAN Limit	Number of IPv6 prefixes that can be delegated for each VLAN
POS Limit	Number of IPv6 prefixes that can be delegated for each POS access interface
Ethernet Limit	Number of IPv6 prefixes that can be delegated for each Ethernet subnet

Table 144: show ipv6 dhcpv6-local limits Output Fields (*continued*)

Field Name	Field Description
Interface	Type of interface and interface specifier. For details about interface types and specifiers, see <i>Interface Types and Specifiers</i> in <i>JunosE Command Reference Guide</i> .
Limit	Number of IPv6 prefixes that can be delegated to the specified interface or subinterface. The limit value configured for the interface type is displayed if no specific limit value is configured for the particular interface.
Count	Number of active IPv6 prefixes delegated on the interface
Denied	Number of IPv6 prefix delegation requests denied during the current denial period. The value is reset to zero (and the denial period restarted) when the number of active delegated IPv6 prefixes no longer exceeds the configured limit.
Total Denied	Total number of IPv6 prefix delegation requests denied on the interface since the interface became active

- Related Documentation**
- [Configuring the DHCPv6 Local Server on page 501](#)
 - `show ipv6 dhcpv6-local limits`

Monitoring the DHCPv6 Local Server to Send Delegated-IPv6-Prefix to SRC Client

Purpose Display the status about DHCPv6 local server to send the Delegated-IPv6-Prefix to the SRC client.

Action To display information whether the DHCPv6 local server is enabled or disabled to send the Delegated-IPv6-Prefix to the SRC client:

```
host1# show ipv6 dhcpv6-local ssc
enabled
```

- Related Documentation**
- [Configuring the DHCPv6 Local Server on page 501](#)
 - `ipv6 dhcpv6-local ssc`
 - `show ipv6 dhcpv6-local ssc`

PART 5

Managing the Subscriber Environment

- [Configuring Subscriber Management on page 605](#)
- [Monitoring Subscriber Management on page 615](#)
- [Configuring Subscriber Interfaces on page 619](#)
- [Monitoring Subscriber Interfaces on page 655](#)

Configuring Subscriber Management

This chapter describes how to set up subscriber management on the E Series router. Subscriber management integrates a variety of router features and enables you to manage your constantly changing subscriber environment without affecting the performance you provide to your customers.

The following sections discuss subscriber management:

- [Understanding Subscriber Management on page 605](#)
- [Subscriber Management Platform Considerations on page 606](#)
- [Subscriber Management Attributes on page 606](#)
- [Subscriber Management Procedure Overview on page 607](#)
- [Configuring Subscriber Management with an External DHCP Server on page 609](#)
- [Subscriber Management Configuration Examples on page 610](#)

Understanding Subscriber Management

The E Series router enables customers to create a unified subscriber management, provisioning, and service delivery environment. The flexibility of the router provides a variety of methods and configurations that enable customers to dynamically provision new subscribers and quickly create new value-added services.

Two major aspects of subscriber management are subscriber provisioning and differentiated service delivery. The E Series router enables you to use both static and dynamic methods to add and delete subscribers. Important subscriber management concepts provided by JunosE subscriber management include:

- Subscriber use of a shared medium
- Multiple subscribers using the same primary interface
- User authentication and accounting
- Differentiated services for individual subscribers

A subscriber management environment can include the following components:

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

- RADIUS server
- Session and Resource Control (SRC) software

You employ the components you need in a variety of configurations, depending on your specific requirements.

- Related Documentation**
- [Subscriber Management Platform Considerations on page 606](#)
 - [Subscriber Management Procedure Overview on page 607](#)

Subscriber Management Platform Considerations

Subscriber management is supported on all E Series routers.

For information about the modules supported on E Series routers:

- See the *ERX Module Guide* for modules supported on ERX7xx models, ERX14xx models, and the ERX310 Broadband Services Router.
- See the *E120 and E320 Module Guide* for modules supported on the E120 and E320 Broadband Services Routers.

- Related Documentation**
- [Understanding Subscriber Management on page 605](#)

Subscriber Management Attributes

E Series routers take advantage of many of the JunosE features to enable you to create the subscriber management environment that best meets your requirements. These features include:

- Authentication—Uses RADIUS to determine whether a user can access a specific service or resource.
- Accounting—Uses RADIUS and policy management to track service usage that can be used for volume-based billing.
- Dynamic address assignment—Uses RADIUS, DHCP, and profiles to dynamically allocate IP addresses to subscribers.
- Dynamic policy management—Uses policy and quality of service (QoS) management to assign and monitor subscriber bandwidth restrictions.
- Security—Uses policy management, source address validation, and media access control (MAC) address validation to grant subscriber access and to enable the use of classification when monitoring subscriber traffic flows.
- Dynamic interfaces—Automatically creates an interface column based on a catalyst packet or event.

- **Marking**—Uses policy management marking to enable differential treatment of specific packets.
- **Policy routing**—Uses policy management routing policies to assign subscriber routes that are based on classification.

Dynamic IP Subscriber Interfaces

You can set up your subscriber management environment to create dynamic IP subscriber interfaces in two situations—when a DHCP event occurs or when a packet is detected.

In the first case, the interface is created when an external DHCP server or the DHCP local server responds to a subscriber request. In the second case, the subscriber interface is created when the router receives a packet (the packet detect feature) with a source IP address that is not in the demultiplexer table. In this second case the primary IP interface must be in autoconfiguration mode.

Subscriber management uses the following process when validating the IP source address of the packet:

- If the address is not valid, no subscriber interface is created. A discard entry is added to the demultiplexer table, and an error message is generated.
- If the address is valid with respect to the address ranges configured on the primary IP interface, subscriber management uses packet information to select the appropriate dynamic subscriber interface profile. The commands corresponding to the profile are then used to create the subscriber interface.

Related Documentation

- [Understanding Subscriber Management on page 605](#)
- [Configuring Subscriber Management with an External DHCP Server on page 609](#)

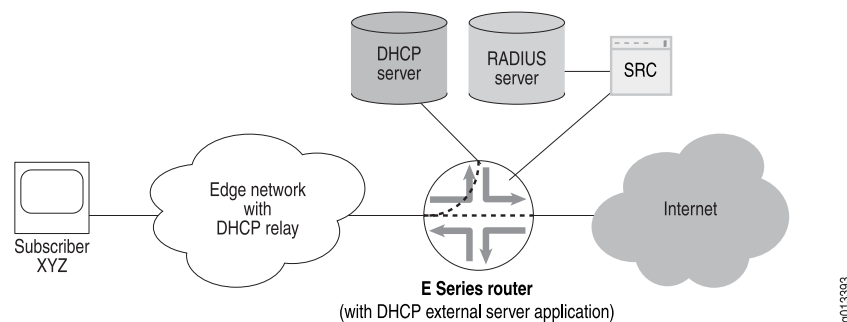
Subscriber Management Procedure Overview

[Figure 15 on page 608](#) shows a subscriber management environment that includes an external DHCP server, a RADIUS server, the SRC software, and the DHCP external server application running on the E Series router.

The E Series router DHCP external server application is used with other JunosE features to provide subscriber management. Using the router's DHCP external server application for subscriber management enables you to take advantage of the following features:

- **Profile assignment**—A dynamic subscriber interface profile is associated with a specific source address by the router's packet detect feature.
- **Dynamic subscriber interface packet detection and inactivity timer**—Subscriber interfaces are dynamically created based on packet information that is identified by the packet detection feature. The inactivity timer determines when a dynamic subscriber interface expires and needs to be deleted.
- **DHCP external server application**—DHCP packets are examined to determine the state of subscribers.

Figure 15: DHCP External Server



In [Figure 15 on page 608](#), the subscriber requests an address from the DHCP server. The E Series router DHCP external server application monitors all DHCP communications between the subscriber and the DHCP server. After the subscriber receives an IP address, the subscriber can access the Internet and use the value-added services provided by the SRC software. The following list describes the various procedures performed in the subscriber management environment:

- Subscriber PC—Requests an IP address from the DHCP server
- E Series router
 - Monitors DHCP traffic between the subscriber and the DHCP server:
 - Identifies the subscriber's IP address, MAC address, giaddr, and client identifier
 - Extracts the lease time, creates a shadow lease, and starts its own lease timer that is associated with the subscriber
 - Determines the subscriber is active when the subscriber sends a packet after receiving an IP address from DHCP. The router then:
 - Processes the subscriber's IP address by using a route map
 - Extracts the dynamic subscriber interface profile (optional)

The router uses the profile to provide authentication, authorization, accounting, and address assignment. RADIUS uses the profile to obtain information for the subscriber's IP interface.
 - Creates the subscriber's dynamic subscriber interface (DSI)
 - If the SRC software is configured, the router also alerts the SRC software that the subscriber's DSI and address exist.
 - The DHCP external server application continues to monitor all traffic between the subscriber and the DHCP server, and periodically resets the shadow lease it originally created when the subscriber first requested an IP address. When the subscriber disconnects, the shadow lease eventually expires, at which time the E Series router performs the following:
 - Deletes the DSI
 - Alerts the SRC software that the DSI has been deleted

- Alerts the SRC software that the subscriber's address has been deleted
- SRC software—Provides enhanced services to the subscriber.

**Related
Documentation**

- [Understanding Subscriber Management on page 605](#)
- [Configuring Subscriber Management with an External DHCP Server on page 609](#)

Configuring Subscriber Management with an External DHCP Server

To configure subscriber management for clients by using an external DHCP server, as in shown in the *DHCP External Server* figure in “[Subscriber Management Procedure Overview](#)” on page 607, use the following procedure on E Series routers:

1. Enable the DHCP external server application.

```
host1(config)#service dhcp-external
```

2. Specify each DHCP server for which to monitor traffic. You can specify a maximum of four DHCP servers.

```
host1(config)#ip dhcp-external server-address 10.10.10.1
```

3. Configure a default policy for subscribers, using a previously configured classifier group.

```
host1(config)#ip policy-list filterAll
host1(config-policy-list)#classifier-group filterGroupA
host1(config-policy-list-classifier-group)#filter
host1(config-policy-list-classifier-group)#exit
host1(config-policy-list)#exit
```

4. Configure a dynamic subscriber interface policy.

```
host1(config)#profile disableUser
host1(config-profile)#ip policy input filterAll
host1(config-profile)#ip policy output filterAll
host1(config-profile)#exit
```

5. Configure a route map.

```
host1(config)#route-map routeMapWest21
host1(config-route-map)#set ip interface-profile disableUser
host1(config-route-map)#exit
```

6. Enable autoconfiguration mode.

```
host1(config)#interface gigabitEthernet 12/0
host1(config-if)#ip address 192.168.1.1 255.255.255.0
host1(config-if)#ip auto-configure ip-subscriber include-primary
host1(config-if)#ip route-map ip-subscriber routeMapWest21
host1(config-if)#ip auto-detect ip-subscriber
host1(config-profile)#exit
```

**Related
Documentation**

- [Subscriber Management Procedure Overview on page 607](#)
- [ip policy](#)

- *ip policy-list*
- *route-map*
- *service dhcp-external*
- *set ip interface-profile*

Subscriber Management Configuration Examples

This section contains examples of creating dynamic usernames and shows the usernames that are generated. The examples all use the following IP policy, interface profile, and route map:

- An IP policy that restricts access.

```
host1(config)#ip policy-list restrictAccess
host1(config-policy-list)#classifier-group *
host1(config-policy-list-classifier-group)#filter
host1(config-policy-list-classifier-group)#exit
host1(config-policy-list)#exit
host1(config)#
```

- An interface profile that references the restrictAccess policy.

```
host1(config)#profile atlInterfaceProfile
host1(config-profile)#ip policy input restrictAccess
host1(config-profile)#ip policy output restrictAccess
host1(config-profile)#exit
host1(config)#
```

- A route map that references the interface profile and the atlServiceProfile service profile.

```
host1(config)#route-map atlRouteMap
host1(config-route-map)#set interface-profile atlInterfaceProfile
host1(config-route-map)#set ip service-profile atlServiceProfile
host1(config-route-map)#exit
host1(config)#
```

The following examples show the configuration of a service profile that enables RADIUS authentication:

- [Username with ATM Circuit Identifier and No Circuit Type on page 610](#)
- [Username with VLAN Circuit Identifier and Circuit Type on page 611](#)
- [Username with MAC Address on page 612](#)

Username with ATM Circuit Identifier and No Circuit Type

This example shows the steps to configure a service profile for a username that includes the ATM circuit identifier, but does not include the circuit type.

```
host1(config)#ip service-profile atlServiceProfile
host1(config-service-profile)#user-prefix xyzcorp.atl
host1(config-service-profile)#domain eastcoast
```

```

host1(config-service-profile)#include hostname
host1(config-service-profile)#include circuit-identifier atm
host1(config-service-profile)#exit
host1(config)#

```

The example generates the following username:

user prefix	circuit identifier	domain
xyzcorp.atl	2.3.32.100.	@eastcoast

The circuit identifier indicates a user at slot 2, port 3, with a virtual path identifier (VPI) of 32 and a virtual channel identifier (VCP) of 100.

Username with VLAN Circuit Identifier and Circuit Type

This example shows the steps to configure a service profile for a username that includes a VLAN circuit identifier and the circuit type.

```

host1(config)#ip service-profile atlServiceProfile
host1(config-service-profile)#user-prefix xyzcorp.atl
host1(config-service-profile)#domain eastcoast
host1(config-service-profile)#include hostname
host1(config-service-profile)#include circuit-identifier vlan prepend-circuit-type
host1(config-service-profile)#exit

```

The example generates the following username:

user prefix	circuit type	circuit identifier	domain
xyzcorp.atl.vlan	1.0.0.	45	@eastcoast

The circuit identifier indicates a user on slot 1, port 0, no stacked vlan, and a vlan ID of 45.

Username with MAC Address

Step-by-Step Procedure This example shows the steps to configure a service profile that generates a username that includes a MAC address.



NOTE: Including a MAC address in a username works only for DHCP users. It does not work for IP subscribers that have statically configured IP addresses.

```
host1(config)#ip service-profile atlServiceProfile
host1(config-service-profile)#user-prefix xyzcorp.atl
host1(config-service-profile)#domain eastcoast
host1(config-service-profile)#include hostname
host1(config-service-profile)#include circuit-identifier vlan
host1(config-service-profile)#include mac-address
host1(config-service-profile)#include dhcp-option 82 agent-circuit-id
host1(config-service-profile)#exit
host1(config)#
```

The example generates the following username, which includes the MAC address:

user prefix	circuit identifier	mac-address	domain
xyzcorp.atl	1.0.0.45	1234.5678.9012	@eastcoast

Related Documentation

- [Understanding Subscriber Management on page 605](#)
- [Subscriber Management Procedure Overview on page 607](#)
- [Configuring Subscriber Management with an External DHCP Server on page 609](#)
- *include circuit-identifier*
- *include dhcp-option 82*
- *include hostname*
- *include ip-address*
- *include mac-address*
- *ip policy*
- *ip policy-list*
- *ip service-profile*
- *profile*
- *route-map*
- *set ip interface-profile*
- *set ip service-profile*

- *user-prefix*

CHAPTER 25

Monitoring Subscriber Management

This chapter describes how to monitor subscriber management on the E Series router.

The following sections describe commands you can use to display status information and statistics for the subscriber management environment:

- [Monitoring IP Service Profiles on page 615](#)
- [Monitoring Active IP Subscribers Created by Subscriber Management on page 616](#)

Monitoring IP Service Profiles

Purpose Display information for all IP service profiles or for a specific profile.

Action To display information about IP service profiles:

```
host1#show ip service-profile
ip service-profile west500
  user-name: finance22
  user-prefix: xyz.bos
  domain: xyzcorp.net
  include virtual-router-name
  include mac-address
  include circuit-identifier atm prepend-circuit-type
  password: 4398aa

ip service-profile at1SerPro9
  user-name: salesCorp
  domain: xyzcorp.net
  include virtual-router-name
  include circuit-identifier vlan
  password: u473qv
```

Meaning [Table 145 on page 615](#) lists the **show ip service-profile** command output field.

Table 145: show ip service-profile Output Fields

Field Name	Field Description
ip service-profile	Name of profile
user-name	Username used to retrieve information from RADIUS for subscriber interfaces

Table 145: show ip service-profile Output Fields (*continued*)

Field Name	Field Description
user-prefix	User prefix used to retrieve information from RADIUS for subscriber interfaces
domain	Domain used to retrieve information from RADIUS for subscriber interfaces
include ip-address	IP address is included in the service profile
include virtual-router-name	Virtual router is included in the service profile
include mac-address	MAC address is included in the service profile
include circuit-identifier	Circuit identifier that is included in the service profile; atm or vlan, and whether the circuit type is prepended
include hostname	Router hostname is included in the service profile
include dhcp-option 82	Suboptions of DHCP option 82 are included in the service profile: agent-circuit-id or agent-remote-id
password	Password used to retrieve information from RADIUS for subscriber interfaces

Related Documentation

- *show ip service-profile*

Monitoring Active IP Subscribers Created by Subscriber Management

Purpose Display information about active IP subscribers that were created by the JunosE Software's subscriber management feature.

Action To display information about subscribers created by subscriber management:

```
host1# show ip-subscriber 2835349506
```

Id	User Name	Ip Address	Virtual Router	Interface
2835349506	user1@isp1.com	192.168.0.1	default	ip192.168.0.1

Id	Login time
2835349506	WED AUG 23 20:46:24 2006

```
host1# show ip-subscriber detail
```

Subscriber List				
Id	User Name	Ip Address	Virtual Router	Interface
2835349506	user1@isp1.com	192.168.0.1	default	ip192.168.0.1 Profile

Id	Login Time		Mac Address	Handle
2835349506	WED AUG 23	20:46:24 2006	3000.0001.9365	13631489
Id	Interface Profile	Service Profile	Option 82	
2835349506	myProfile	profile22	FastEthernet 3/1	

Meaning [Table 146 on page 617](#) lists the **show ip-subscriber** command output field.

Table 146: show ip-subscriber Output Fields

Field Name	Field Description
Id	ID of the subscriber
Username	Username used to retrieve information from RADIUS for the subscriber interface
Ip Address	IP address of the subscriber interface
Virtual Router	Name of the virtual router on which the subscriber interface is configured
Interface	Name of subscriber interface; ip indicates that subscriber manager created this interface
Login Time	Day, date, and time that the subscriber logged in
Mac Address	MAC address of the subscriber
Profile Handle	AAA profile handle
Interface Profile	Interface profile name used to configure the subscriber interface
Service Profile	IP service profile name used by subscriber management to authorize and configure the subscriber interface with AAA
Option 82	DHCP relay agent information (option 82) circuit identifier that describes the physical interface location associated with the subscriber

Related Documentation

- *show ip-subscriber*

CHAPTER 26

Configuring Subscriber Interfaces

This chapter describes how to configure static and dynamic subscriber interfaces for remote access to the E Series router. This chapter contains the following sections:

- [Subscriber Interfaces Overview on page 619](#)
- [Dynamic Interfaces and Dynamic Subscriber Interfaces Overview on page 620](#)
- [Subscriber Interfaces Platform Considerations on page 622](#)
- [Subscriber Interfaces References on page 623](#)
- [Characteristics of Subscriber Interfaces on page 623](#)
- [Applications for Subscriber Interfaces on page 626](#)
- [Dynamic Creation of Subscriber Interfaces Overview on page 628](#)
- [Dynamic Subscriber Interfaces Using DHCP Events Overview on page 629](#)
- [Dynamic Subscriber Interfaces Using Packet Detection Overview on page 631](#)
- [Inheritance of MAC Address Validation State for Dynamic Subscriber Interfaces Overview on page 632](#)
- [How MAC Address Validation State Inheritance Works on page 633](#)
- [Example: Configuring Static Subscriber Interfaces on page 635](#)
- [Example: Configuring Dynamic Subscriber Interfaces on page 642](#)
- [Example: Configuring Dynamic Subscriber Interfaces Using Loopback Interfaces on page 650](#)

Subscriber Interfaces Overview

You can configure E Series routers to create subscriber interfaces statically or dynamically.

The following list shows the underlying (layer 2) interfaces on which you can currently configure each type of subscriber interface.

- Static subscriber interfaces
 - Bridged Ethernet over ATM (with and without VLANs)
 - Fast Ethernet (with and without VLANs)
 - Gigabit Ethernet (with and without VLANs)

- 10-Gigabit Ethernet (with and without VLANs)
- IP over ATM
- POS
- Generic Routing Encapsulation (GRE) tunnels
- Dynamic subscriber interfaces
 - Bridged Ethernet over ATM (with and without VLANs)
 - Fast Ethernet (with and without VLANs)
 - Gigabit Ethernet (with and without VLANs)
 - 10-Gigabit Ethernet (with and without VLANs)
 - GRE tunnels

For information about platform support for subscriber interfaces, see [“Subscriber Interfaces Platform Considerations” on page 622](#).

**Related
Documentation**


- [Dynamic Interfaces and Dynamic Subscriber Interfaces Overview on page 620](#)
- [Characteristics of Subscriber Interfaces on page 623](#)
- [Dynamic Creation of Subscriber Interfaces Overview on page 628](#)
- [Applications for Subscriber Interfaces on page 626](#)
- [Subscriber Interfaces References on page 623](#)

Dynamic Interfaces and Dynamic Subscriber Interfaces Overview

Dynamic interfaces are created automatically and transparently in response to external events. For example, the router creates dynamic interfaces when a lower-layer link such as an ATM or VLAN receives data. The layers of a dynamic interface are created based on the packets received on the link and can be configured using profiles, RADIUS, or a combination of the two. Dynamic interfaces are used to terminate Broadband Residential Access Server (B-RAS) access such as: Point-to-Point Protocol over Ethernet (PPPoE), Point-to-Point Protocol over ATM (PPPoA), and Point-to-Point Protocol over Ethernet over ATM (PPPoEoA). A PPP session acts as logical separation between one subscriber session and the next. Multiple services using policies and QoS can be applied to the IP interface that is associated with the PPP session.

An example of a dynamic interface configuration is a PPPoE session running on top of a Gigabit Ethernet VLAN interface. [Figure 16 on page 621](#) shows an example of the dynamic interface stack.


Figure 16: Example of a Dynamic Interface Stack



You can configure the lower layers of the stack (GE physical interface and VLAN major interface) either dynamically or statically, and dynamically configure the upper layers (VLAN subinterface, PPPoE, and IP). An interface is considered dynamic if at least one of the layers in the interface stack is configured dynamically.

The router creates dynamic subscriber interfaces (DSIs) on demand, in response to external events, such as when a Dynamic Host Configuration Protocol (DHCP) event occurs or when the router detects a packet. DSIs function in a manner similar to dynamic interfaces. However, DSIs have a more specific application than dynamic interfaces. You use DSIs when there are no PPPoE, PPPoA, or PPPoEoA sessions to provide separation between layers and when subscriber management is required. For example, on an Ethernet VLAN, multiple subscribers can enter the network from a Wi-Fi hotspot, as shown in [Figure 17 on page 621](#):

Figure 17: Example of a Dynamic Subscriber Interface



In [Figure 17 on page 621](#), multiple subscribers share the same broadcast segment. Each subscriber is identified by an individual IP address or a group of subscribers can be identified with an IP network. When each subscriber is identified by an individual IP address, a dynamic subscriber interface is created for each subscriber. You can manage a group of subscribers identified with an IP network, on a single DSI. You can also manage a group of subscribers using a static subscriber interface (SSI). However, you must

manually configure the SSI and you cannot use the same dynamic profiles and RADIUS that DSIs use.

Subscribers can be connected to a single broadcast segment without using dynamic or static subscriber interfaces. This configuration is useful when subscriber management is not required. Subscriber management usually refers to (but is not limited to) tailoring IP policies and QoS profiles to a specific address or a very small group of addresses. For detailed information about the uses for Dynamic Subscriber interfaces, see [“Example: Configuring Dynamic Subscriber Interfaces” on page 642](#).

**Related
Documentation**

- [Subscriber Interfaces Overview on page 619](#)
- [Characteristics of Subscriber Interfaces on page 623](#)
- [Dynamic Creation of Subscriber Interfaces Overview on page 628](#)
- [Applications for Subscriber Interfaces on page 626](#)
- [Subscriber Interfaces References on page 623](#)

Subscriber Interfaces Platform Considerations

For information about modules that support subscriber interfaces on the ERX7xx models, ERX14xx models, and the ERX310 Broadband Services Router:

- See *ERX Module Guide, Table 1, ERX Module Combinations* for detailed module specifications.
- See *ERX Module Guide, Appendix A, Module Protocol Support* for information about the modules that support subscriber interfaces.

For information about modules that support subscriber interfaces on the E120 and E320 Broadband Services Routers:

- See *E120 and E320 Module Guide, Table 1, Modules and IOAs* for detailed module specifications.
- See *E120 and E320 Module Guide, Appendix A, IOA Protocol Support* for information about the modules that support subscriber interfaces.

Interface Specifiers

The configuration task examples in this chapter use the *slot/port* format to specify an interface. However, the interface specifier format that you use depends on the router that you are using.

For ERX7xx models, ERX14xx models, and ERX310 routers, use the *slot/port* format. For example, the following command specifies a Gigabit Ethernet interface on slot 0, port 1 of an ERX7xx model, ERX14xx model, or ERX310 Broadband Services Router.

```
host1(config)#interface gigabitEthernet 0/1
```

For E120 and E320 Routers, use the *slot/adaptor/port* format, which includes an identifier for the bay in which the I/O adapter (IOA) resides. In the software, adaptor 0 identifies

the right IOA bay (E120 router) and the upper IOA bay (E320 router); adapter 1 identifies the left IOA bay (E120 router) and the lower IOA bay (E320 router). For example, the following command specifies a Gigabit Ethernet interface on slot 5, adapter 0, and port 0 of an E320 router.

```
host1(config)#interface gigabitEthernet 5/0/0
```

For more information about supported interface types and specifiers on E Series routers, see *Interface Types and Specifiers* in *JunosE Command Reference Guide*.

**Related
Documentation**

- [Subscriber Interfaces Overview on page 619](#)
- [Dynamic Interfaces and Dynamic Subscriber Interfaces Overview on page 620](#)
- [Characteristics of Subscriber Interfaces on page 623](#)
- [Dynamic Creation of Subscriber Interfaces Overview on page 628](#)
- [Applications for Subscriber Interfaces on page 626](#)

Subscriber Interfaces References

For more information about the DHCP local server and DHCP external server, which are used in dynamic creation of subscriber interfaces, consult the following resources:

- [“DHCP Overview Information” on page 461](#)
- RFC 2131—Dynamic Host Configuration Protocol (March 1997)

**Related
Documentation**

- [Subscriber Interfaces Overview on page 619](#)
- [Dynamic Interfaces and Dynamic Subscriber Interfaces Overview on page 620](#)
- [Characteristics of Subscriber Interfaces on page 623](#)
- [Dynamic Creation of Subscriber Interfaces Overview on page 628](#)
- [Applications for Subscriber Interfaces on page 626](#)

Characteristics of Subscriber Interfaces

The following sections describe the characteristics of subscriber interfaces:

- [Relationship to Shared IP Interfaces on page 624](#)
- [Relationship to Primary IP Interfaces on page 624](#)
- [Ethernet Interfaces and VLANs on page 625](#)
- [Moving Interfaces on page 625](#)
- [Preventing IP Spoofing on page 625](#)
- [Routing Protocols on page 625](#)
- [Policies and QoS on page 625](#)

Relationship to Shared IP Interfaces

A subscriber interface is an extension of a *shared IP interface*. A shared IP interface is one of a group of IP interfaces that use the same layer 2 interface.

Shared IP interfaces are unidirectional—they can transmit but not receive traffic. In contrast, subscriber interfaces are bidirectional—they can both receive and transmit traffic.

For details about shared IP interfaces, see the *Shared IP Interfaces* section in the *JunosE IP, IPv6, and IGP Configuration Guide*.

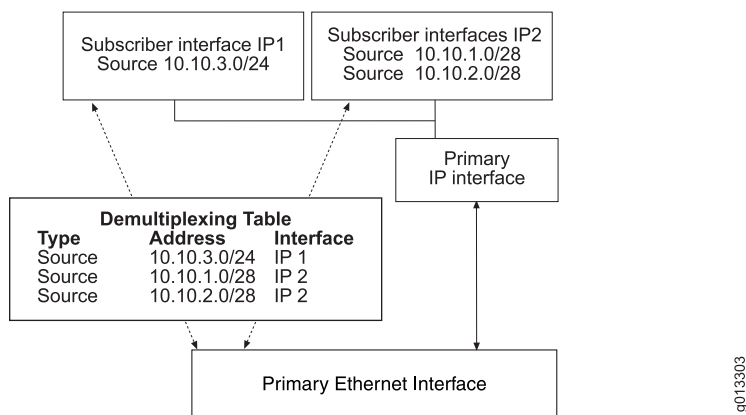
Relationship to Primary IP Interfaces

A subscriber interface operates only with a *primary IP interface*—a normal IP interface on a supported layer 2 interface, such as Ethernet. You create a primary interface by assigning an IP address to the Ethernet interface. Although you can configure a subscriber interface directly on an Ethernet interface, the subscriber interface does not operate until you assign an IP address to the Ethernet interface.

To configure a subscriber interface you must associate either a source address or a destination address with the interface. The router receives packets on a subscriber interface after demultiplexing the packet according to the specified source address or destination address. You can associate multiple source addresses or multiple destination addresses with a subscriber interface. However, a single primary interface and its associated subscriber interfaces can only demultiplex source addresses or destination addresses at any given time.

For example, [Figure 18 on page 624](#) illustrates the relationship between subscriber interfaces, an associated primary IP interface, and an associated Ethernet interface.

Figure 18: Subscriber Interfaces over Ethernet



When the router receives traffic on a primary interface, the primary interface performs a lookup in its demultiplexing table. If the result of the lookup is a subscriber interface, the traffic is received on the associated subscriber interface.



NOTE: You can use the `set dhcp relay giaddr-selects-interface` command to specify that the primary interface is identified by information in the `giaddr` field of DHCP ACK messages. By default, the router identifies the primary interface based on the interface used by the DHCP-destined packets. See the *Using the Giaddr to Identify the Primary Interface for Dynamic Subscriber Interfaces* section in “DHCP Relay and BOOTP Relay Overview” on page 512.

Ethernet Interfaces and VLANs

In the absence of VLANs, Ethernet does not have a demultiplexing layer. A subscriber interface adds a demultiplexing layer for an Ethernet interface that is configured without VLANs. Using subscriber interfaces, the router can demultiplex or separate the traffic associated with different subscribers.

You can configure subscriber interfaces with VLANs. If you do so, the E Series router demultiplexes packets by using first the VLAN and then the subscriber interface.

Moving Interfaces

A shared IP interface that has associated subscriber demultiplexing attributes retains these attributes when it moves.

For details about moving shared IP interfaces, see the *Moving IP Interfaces* section in the *JunosE IP, IPv6, and IGP Configuration Guide*.

Preventing IP Spoofing

You can prevent IP spoofing on subscriber interfaces by using media access control (MAC) address validation.

For information about configuring MAC address validation, see the *MAC Address Validation* section in the *JunosE IP, IPv6, and IGP Configuration Guide*.

For information about the relationship between the MAC address validation state and dynamically created subscriber interfaces, see the *Inheritance of MAC Address Validation State for Dynamic Subscriber Interfaces* section in “DHCP Relay and BOOTP Relay Overview” on page 512.

Routing Protocols

You configure unicast routing protocols on subscriber interfaces in the same way that you configure routing protocols on primary IP interfaces, provided that you configure them to use unicast addressing when communicating with a peer. You can also enable multicast routing protocols such as IGMP on subscriber interfaces; however, we do not recommend this type of configuration.

Policies and QoS

You can configure policies, such as rate limiting and filtering, and quality of service (QoS) for subscriber interfaces in the same way that you do for primary IP interfaces. For more

information, see the *JunosE Policy Management Configuration Guide* and the *JunosE Quality of Service Configuration Guide*.

**Related
Documentation**

- [Subscriber Interfaces Overview on page 619](#)
- [Dynamic Interfaces and Dynamic Subscriber Interfaces Overview on page 620](#)
- [Dynamic Creation of Subscriber Interfaces Overview on page 628](#)
- [Applications for Subscriber Interfaces on page 626](#)
- [Example: Configuring Static Subscriber Interfaces on page 635](#)
- [Example: Configuring Dynamic Subscriber Interfaces on page 642](#)

Applications for Subscriber Interfaces

In a cable modem network, service providers can use subscriber interfaces to:


- Direct traffic toward special local content in the network
- Differentiate traffic for virtual private networks (VPNs)

Directing Traffic Toward Special Local Content

[Figure 19 on page 627](#) shows an example of a cable modem network. Multiple cable modem termination systems (CMTSs) connect to multiple shared media access LANs. Many subscribers connect to each LAN.

In this example, the service provider uses subscriber interfaces to direct traffic toward special local content on the network: a voice over Internet Protocol (VoIP) service on network 10.11.0.0/16, or a local gaming service on network 10.12.0.0/16. Rate limits and policies on the subscriber interface customize the service level for the associated service. In this application, the E Series router is the first-hop router for the subscribers, and the subscriber interfaces demultiplex traffic based on the destination address.

Figure 19: Subscriber Interfaces in a Cable Modem Network




For instructions on configuring the application shown in [Figure 19 on page 627](#), see [“Example: Using a Destination Address to Demultiplex Traffic” on page 635](#).

Differentiating Traffic for VPNs

Similarly, service providers can use subscriber interfaces to differentiate traffic for VPNs. [Figure 20 on page 628](#) shows an example of this application.

Customers on subnet A need to connect to VPN A, and customers on subnet B need to connect to VPN B. The E Series router connects to VPN A through virtual router A and to VPN B through virtual router B. Using two subscriber interfaces on the same primary interface (one on virtual router B and one on virtual router A), the E Series router can separate the traffic from subnets A and B. Because the E Series router is forwarding traffic in this application, the shared IP interface should demultiplex the traffic by using a source address.

Figure 20: Associating Subnets with a VPN Using Subscriber Interfaces



For instructions on configuring the application shown in Figure 20 on page 628, see “Example: Using a Destination Address to Demultiplex Traffic” on page 635.

Related Documentation

- [Subscriber Interfaces Overview on page 619](#)
- [Dynamic Interfaces and Dynamic Subscriber Interfaces Overview on page 620](#)
- [Characteristics of Subscriber Interfaces on page 623](#)
- [Dynamic Creation of Subscriber Interfaces Overview on page 628](#)
- [Example: Configuring Static Subscriber Interfaces on page 635](#)
- [Example: Configuring Dynamic Subscriber Interfaces on page 642](#)

Dynamic Creation of Subscriber Interfaces Overview

As an alternative to creating static subscriber interfaces, you can configure E Series routers to create subscriber interfaces dynamically.

When you create a static subscriber interface, as described in “Example: Configuring Static Subscriber Interfaces” on page 635, each layer in the interface stack is created through an existing configuration mechanism such as command-line interface (CLI) or Simple Network Management Protocol (SNMP).

By contrast, the router creates dynamic subscriber interfaces on demand, in response to an external event. Two types of external events can cause dynamic creation of subscriber interfaces: when a Dynamic Host Configuration Protocol (DHCP) event occurs or when the router detects a packet.

The following topics describe the two methods of dynamic creation of subscriber interface columns:

**Related
Documentation**

- [Dynamic Subscriber Interfaces Using DHCP Events Overview on page 629](#)
- [Dynamic Subscriber Interfaces Using Packet Detection Overview on page 631](#)
- [Subscriber Interfaces Overview on page 619](#)
- [Dynamic Interfaces and Dynamic Subscriber Interfaces Overview on page 620](#)
- [Characteristics of Subscriber Interfaces on page 623](#)
- [Applications for Subscriber Interfaces on page 626](#)
- [Example: Configuring Dynamic Subscriber Interfaces on page 642](#)

Dynamic Subscriber Interfaces Using DHCP Events Overview

The DHCP event that triggers dynamic creation of subscriber interfaces occurs when either a local DHCP server or external DHCP server assigns an IP address to a subscriber that has issued a DHCP request. After the DHCP server assigns the IP address and the router creates the associated dynamic subscriber interface, the subscriber can access required network services.

DHCP Local Server and Address Allocation

You can configure the DHCP local server to operate in either equal-access mode or standalone mode.

In standalone mode, the DHCP local server provides a basic DHCP service. The server receives a client request for an IP address and immediately allocates the subscriber an IP address from one of the local address pools.

In equal-access mode, the DHCP local server works with Juniper Networks Session and Resource Control (SRC) software and the authorization, accounting, and address assignment utility to provide an advanced subscriber configuration and management service. After the subscriber is authenticated through RADIUS, the DHCP server assigns the subscriber an IP address with a long lease time. This assignment of an IP address triggers the creation of dynamic subscriber interfaces.

For more information about the DHCP servers and the SRC software, see the following chapters:

- [“DHCP Overview Information” on page 461](#)
- *SRC-PE Getting Started Guide, SRC Product Overview*

DHCP External Server and Address Allocation

With DHCP external server, all communication between the subscriber and the DHCP server is monitored by the E Series router. The subscriber requests an address from the DHCP server through the E Series router. After the subscriber receives an IP address, the subscriber can access the Internet and use the value-added services provided by the E Series router and by the SRC software. The edge network must be using a DHCP relay function.

The services provided by integrating the E Series router's DHCP external server application with SRC software are similar to those provided when the DHCP local server is integrated with SRC software. For more information, see *SRC-PE Getting Started Guide*, *SRC Product Overview*.

DHCP Relay Configuration

When you are configuring dynamic subscriber interface support, and you configure DHCP relay in the same virtual router as the dynamic subscriber interfaces, you must use the **set dhcp relay inhibit-access-route-creation** command to ensure that DHCP relay does not install access internal routes. Otherwise, DHCP relay will overwrite the access internal routes that are originally created for the subscriber interface.

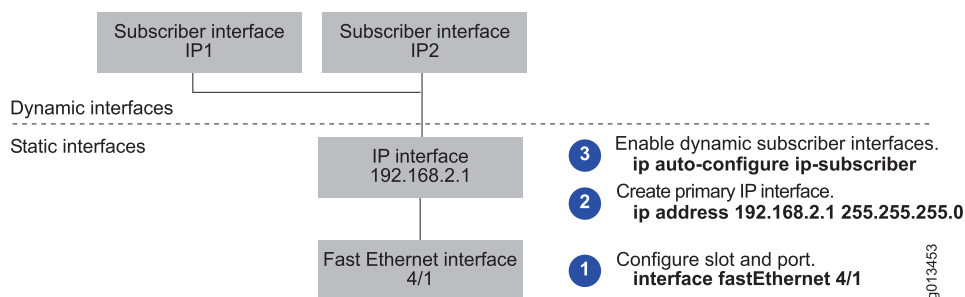
Supported Configurations

E Series routers currently support dynamic creation of subscriber interfaces with DHCP servers in the following configurations:

- IP over Ethernet
- IP over VLAN over Ethernet
- IP over bridged Ethernet over ATM

For example, [Figure 21 on page 630](#) shows the interface stacking in an IP over Ethernet dynamic subscriber interface configuration. The illustration indicates which layers in the stack are static and dynamic, and identifies the CLI commands typically used to create the configuration.

Figure 21: IP over Ethernet Dynamic Subscriber Interface Configuration



As shown in [Figure 21 on page 630](#), issuing the **ip auto-configure ip-subscriber** command configures the primary IP interface to enable dynamic creation of subscriber interfaces.

However, the router does not actually create the dynamic subscriber interface until the DHCP server assigns an IP address to the associated subscriber.

To configure each supported configuration, see [“Example: Configuring Dynamic Subscriber Interfaces” on page 642](#).

**Related
Documentation**

- [Subscriber Interfaces Overview on page 619](#)
- [Dynamic Interfaces and Dynamic Subscriber Interfaces Overview on page 620](#)
- [Characteristics of Subscriber Interfaces on page 623](#)
- [Applications for Subscriber Interfaces on page 626](#)
- [Dynamic Subscriber Interfaces Using Packet Detection Overview on page 631](#)
- [Example: Configuring Dynamic Subscriber Interfaces on page 642](#)

Dynamic Subscriber Interfaces Using Packet Detection Overview

For GRE tunnel interfaces, the event that triggers dynamic creation of subscriber interfaces occurs when the router receives a packet with a source IP address that is not in the demultiplexer table. In this case, the primary IP interface must be in autoconfiguration mode.

Packet detection is the only method of dynamically creating subscriber interfaces on GRE tunnel interfaces; you cannot use DHCP local server or DHCP external server.

Issuing the **ip auto-configure ip-subscriber** command configures the primary IP address to enable dynamic configuration of subscriber interfaces. Unlike DHCP configurations, the router creates the dynamic subscriber interface when it receives the first packet that contains the subscriber's IP address as the source address.

In addition, a dynamic subscriber interface becomes inactive after a period of time in which the router receives no packets that contain the subscriber's IP address as the source address. You can configure the period of time by issuing the **ip inactivity-timer** command.

To configure dynamic creation of subscriber interfaces on GRE tunnel interfaces, see [“Example: Configuring Dynamic Subscriber Interfaces over GRE Tunnels” on page 648](#).

Designating Traffic for the Primary IP Interface

When dynamic creation of subscriber interfaces is enabled on the primary IP interface (by means of the **ip auto-configure ip-subscriber** command), you can use the **ip source-prefix** command to specify the source address of traffic that is destined for the primary IP interface instead of the subscriber interface. If the DHCP server (for DHCP server configurations) or the router (for packet detection configurations) then assigns a subscriber an IP address matching this source prefix, the router does not create a dynamic subscriber interface for that address.

Using Framed Routes

You can use the **ip use-framed-routes ip-subscriber** command to enable a primary IP interface to use framed routes as source IP addresses when creating dynamic subscriber interfaces. The framed routes are applied to the dynamic subscriber interface during configuration so traffic from the subsets can traverse the interface. By applying framed routes in this fashion, you can extend the per-subscriber interface management to any subnetworks behind the dynamic subscriber interface. RADIUS includes the Framed-Route attribute [22] in Access-Accept messages to specify the route in the following format:

Framed-Route = *ipAddress/mask nextHop*

Related Documentation

- [Subscriber Interfaces Overview on page 619](#)
- [Dynamic Interfaces and Dynamic Subscriber Interfaces Overview on page 620](#)
- [Characteristics of Subscriber Interfaces on page 623](#)
- [Dynamic Creation of Subscriber Interfaces Overview on page 628](#)
- [Dynamic Subscriber Interfaces Using DHCP Events Overview on page 629](#)
- [Applications for Subscriber Interfaces on page 626](#)
- [Example: Configuring Dynamic Subscriber Interfaces on page 642](#)

Inheritance of MAC Address Validation State for Dynamic Subscriber Interfaces Overview

A dynamic IP subscriber interface inherits the MAC address validation state (enabled or disabled) configured for its parent static primary IP interface.

MAC address validation binds a MAC source address for an interface to a given IP source address. When the IP-MAC binding is established, the router forwards ingress packets on the interface when the packet's MAC source address and IP source address match, and drops ingress packets when the packet's MAC source address and IP source address do not match. MAC address validation thereby prevents spoofing on IP-based Ethernet interfaces, and is very useful in subscriber management applications.

When MAC address validation is enabled on an interface, the router checks the entry in the MAC validation table that corresponds to the IP source address of an incoming packet. The MAC source address of the packet must match the MAC source address of the table entry for the router to forward the packet.

Related Documentation

- [Subscriber Interfaces Overview on page 619](#)
- [Dynamic Interfaces and Dynamic Subscriber Interfaces Overview on page 620](#)
- [Characteristics of Subscriber Interfaces on page 623](#)
- [Dynamic Creation of Subscriber Interfaces Overview on page 628](#)
- [Applications for Subscriber Interfaces on page 626](#)
- [How MAC Address Validation State Inheritance Works on page 633](#)

- [Example: Configuring Dynamic Subscriber Interfaces on page 642](#)

How MAC Address Validation State Inheritance Works

To enable MAC address validation for the static primary IP interface, you must use the existing **ip mac-validate** command with either the **strict** keyword or the **loose** keyword. The **strict** keyword prevents transmission of IP packets that do not reside in the MAC validation table. Use the **strict** keyword to forward packets only when both the IP source address and the MAC source address match one of the IP-MAC address pair entries in the table. When the MAC address in the table does not match the MAC source address, or when IP source address of the incoming packet does not match any of the IP addresses in the validation table, the packet is dropped.

The **loose** keyword forwards packets when both the IP source address and the MAC source address match one of the IP-MAC address pair entries in the MAC validation table. When the IP source address matches one of the IP source addresses in the table, but the MAC address of the incoming packet does not match the MAC address of the entry in the table, the packet is dropped. However, when the IP source address of the incoming packet does not match any of the IP addresses in the table, the packet is forwarded. This is the default setting.



NOTE: When a DHCP discover or a DHCP request packet arrives from a requesting client to the router that functions as the DHCP server or the delegating router on an interface, and if you configured either strict or loose mode of MAC address validation on that interface, the DHCP discover or request packets are processed correctly and are not dropped.

When a dynamic IP subscriber interface is created with the MAC address validation state inherited from the static primary IP interface, an entry for the MAC source address is installed in the MAC validation table when MAC address validation is enabled (either loose or strict) on the static primary IP interface. For each packet received on this interface, the router compares the packet's MAC source address to the value in the MAC validation table. If these values match, the router forwards the packet; otherwise, the packet is discarded.

In addition, creation of the dynamic IP subscriber interface adds a static MAC address validation entry in the router's Address Resolution Protocol (ARP) table. This occurs regardless of whether you configure MAC address validation on the static primary IP interface with the **ip mac-validate strict** command or the **ip mac-validate loose** command.

Configuration of MAC Address Validation State Inheritance

No special configuration is required to enable inheritance of the MAC address validation state on dynamic IP subscriber interfaces; this occurs automatically provided that MAC address validation is properly enabled on the parent static primary IP interface with the **ip mac-validate** command. If MAC address validation is disabled on the static primary

IP interface, the dynamic subscriber interface inherits the disabled state for MAC address validation.

Keep the following guidelines in mind for using dynamic IP subscriber interfaces that inherit the MAC address validation state from their parent static primary IP interface:

- A dynamic subscriber interface inherits the MAC address validation state of its static primary IP interface only when the dynamic subscriber interface is created.
- You cannot change the MAC address validation state inherited by a dynamic subscriber interface from its static primary IP interface.
- Changing the MAC address validation state of a static primary IP interface does not affect the MAC address validation state of dynamic subscriber interfaces already created from this primary IP interface. Any dynamic subscriber interfaces created from this primary IP interface after you change the MAC address validation state inherit the new MAC validation state.
- When you configure a dynamic subscriber interface with one or more framed routes (subnets), we recommend that you use the **ip mac-validate loose** command to configure MAC address validation for the static primary IP interface. Using the **loose** keyword, which is the default, prevents the router from discarding packets with an IP source address from a subnet.
- Because enabling MAC address validation on an IP interface creates a static MAC address validation entry in the router's ARP table, be sure to observe the system limit for the maximum number of dynamic ARP table entries supported per line module. See the Link Layer Maximums tables in *Appendix A, System Maximums*, of the *Release Notes* corresponding to your software release for information about the maximum number of dynamic ARP entries that the router supports. Currently, this limit is set to 32,768 dynamic ARP entries for all E Series modules that support Ethernet interfaces.

Verification of MAC Address Validation State Inheritance

To verify inheritance of the MAC address validation state on a dynamic subscriber interface, you can use the **show ip mac-validate interface** command and the **show arp** command.

The following sample output from the **show ip mac-validate interface** command displays the MAC address validation state (strict) inherited by the dynamic subscriber interface ip74.39.64.3 from its parent static primary IP interface.

```
host1#show ip mac-validate interface ip74.39.64.3
ip74.39.64.3: Strict
```

Address	Hardware Addr
74.39.64.3	0090.1a40.f4f6

Building on this example, the following sample output from the **show arp** command displays a static MAC address validation entry (74.39.64.3) in the ARP table for the dynamic subscriber interface when it is created with the MAC address validation state inherited from its parent static primary IP interface. The asterisk (*) indicates that the

ARP entry was added as the result of issuing an **arp validate** command rather than an **arp** command.

```
host1#show arp
      Address      Age      Hardware Addr      Interface
10.13.10.1        21600      0090.6939.751b      FastEthernet6/0
74.39.64.3         -          0090.1a40.f4f6      ip74.39.64.3 *
192.168.1.2       20700      0090.1a40.280d      FastEthernet8/2
```

Related Documentation

- [Subscriber Interfaces Overview on page 619](#)
- [Dynamic Interfaces and Dynamic Subscriber Interfaces Overview on page 620](#)
- [Characteristics of Subscriber Interfaces on page 623](#)
- [Dynamic Creation of Subscriber Interfaces Overview on page 628](#)
- [Applications for Subscriber Interfaces on page 626](#)
- [Inheritance of MAC Address Validation State for Dynamic Subscriber Interfaces Overview on page 632](#)
- [Example: Configuring Static Subscriber Interfaces on page 635](#)
- [Example: Configuring Dynamic Subscriber Interfaces on page 642](#)

Example: Configuring Static Subscriber Interfaces

You can configure static subscriber interfaces on ATM, Fast Ethernet, Gigabit Ethernet, 10-Gigabit Ethernet, or POS layer 2 interfaces.

The examples in this topic show how to configure static subscriber interfaces on a Fast Ethernet interface, but the steps for configuring static subscriber interfaces over other supported layer 2 interface types are similar.

- [Example: Using a Destination Address to Demultiplex Traffic on page 635](#)
- [Example: Using a Source Address to Demultiplex Traffic on page 638](#)

Example: Using a Destination Address to Demultiplex Traffic

This example shows how you can use static subscriber interfaces to direct traffic toward special local content on the network, based on the traffic's destination address.

- [Requirements on page 635](#)
- [Overview on page 636](#)
- [Using a Destination Address to Demultiplex Traffic on page 636](#)

Requirements

This example uses the following software and hardware components:


- JunosE Release 7.1.0 or higher-numbered releases
- E Series router (ERX7xx models, ERX14xx models, the ERX310 router, the E120 router, or the E320 router)

- ASIC-based line modules that support Fast Ethernet or Gigabit Ethernet

Overview

The example in [Figure 22 on page 636](#) shows how you can use static subscriber interfaces to direct traffic toward special local content on the network, based on the traffic's destination address. In this application, a local VoIP service is on network 10.11.0.0/16, and a local gaming service is on network 10.12.0.0/16.

Figure 22: Subscriber Interfaces Using a Destination Address to Demultiplex Traffic



Using a Destination Address to Demultiplex Traffic

To configure the static subscriber interfaces shown in [Figure 22 on page 636](#), perform the following steps:

1. Configure a primary IP interface on a supported layer 2 interface.
 - a. Create a layer 2 interface.


```
host1(config)#interface fastEthernet 3/1
```
 - b. Create a primary IP interface.


```
host1(config-if)#ip address 10.1.1.1 255.0.0.0
```

- c. Configure the primary interface to use a destination address to demultiplex traffic for the subscriber interface. (By default, a source address is used to demultiplex traffic.)

```
host1(config-if)#ip demux-type da-prefix
```

Use the **no** version to restore the default situation in which the router uses a source address to demultiplex traffic.

- d. Exit Interface Configuration mode.

```
host1(config-if)#exit
```

2. Configure subscriber interface IP1.

- a. Create the shared IP interface.

```
host1(config)#interface ip ip1
```

Use the specified name to refer to the shared IP interface; you cannot use the layer 2 interface to refer to the shared IP interface, because the shared interface can be moved.

Use the **no** version to delete the IP interface.

- b. Associate the shared IP interface with the layer 2 interface by using one of the following methods:

- Static

```
host1(config-if)#ip share-interface fastEthernet 3/1
```

If you issue this command on a shared IP interface, you cannot issue the **ip share-nexthop** command for the interface.

After creating the shared IP interface, you can configure it as you do any other IP interface. The shared interface is operationally up when the layer 2 interface is operationally up and IP is properly configured. You can create operational shared IP interfaces in the absence of a primary IP interface.

Use the **no** version to remove the association between the layer 2 interface and the shared IP interface. You can delete shared and primary IP interfaces independently.

- Dynamic

```
host1:vr-a:vrf-1(config-if)#ip share-nexthop 10.1.1.2
```

You can use this command to specify that the shared IP interface dynamically tracks a next hop. If the next hop changes, the shared IP interface moves to the new layer 2 interface associated with the IP interface toward the new next hop.

If you issue this command on a shared IP interface, you cannot issue the **ip share-interface** command for the interface. If you issue this command on a shared IP interface, the shared interface cannot dynamically track the next hop for the specified destination if the next-hop IP address is resolvable over MPLS. If you specify a virtual router, the command fails if the VR does not already exist. If you do not specify a VR, the current VR is assumed.

After creating the shared IP interface, you can configure it as you do any other IP interface. The shared interface is operationally up when the layer 2 interface associated with the specified next hop is operationally up and IP is properly configured.

Use the **no** version to halt tracking of the next hop.

- c. To fully configure the shared interface, assign an address or make it unnumbered.

```
host1(config-if)#ip unnumbered loopback 0
```

- d. Specify the destination addresses for the subscriber interface to use to demultiplex traffic.

```
host1(config-if)#ip destination-prefix 10.11.0.0 255.255.0.0
```

On the ERX1440 Broadband Services Router or the E320 router, you can configure up to 1024 subnets for static subscriber interfaces per primary IP interface when each subnet has a variable network mask that is less than /32. The number of subnets identifying a single route (/32) is still limited by the global maximum of 16,000 hosts per line module.

Use the **no** version to remove the association between the interface and the specified IP destination address and mask.

- e. Exit Interface Configuration mode.

```
host1(config-if)#exit
```

3. Repeat Step 2 to configure subscriber interface IP2.

```
host1(config)#interface ip ip2
host1(config-if)#ip share-interface fastEthernet 3/1
host1(config-if)#ip unnumbered loopback 0
host1(config-if)#ip destination-prefix 10.12.0.0 255.255.0.0
```

Example: Using a Source Address to Demultiplex Traffic

This example illustrates how you can use static subscriber interfaces to differentiate traffic for VPN access, based on the traffic's source address.

- [Requirements on page 638](#)
- [Overview on page 639](#)
- [Using a Source Address to Demultiplex Traffic on page 639](#)

Requirements

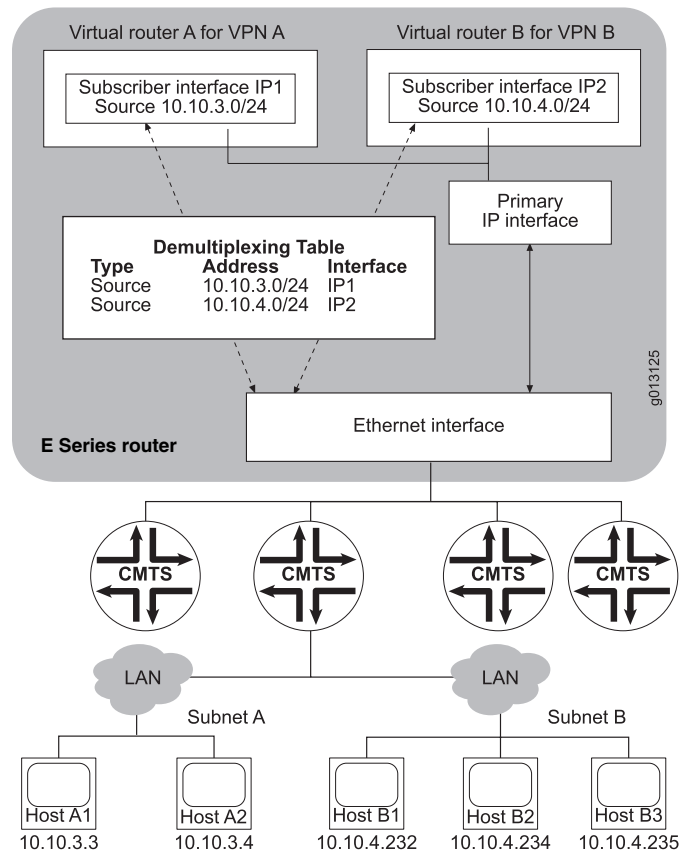
This example uses the following software and hardware components:

- JunosE Release 7.1.0 or higher-numbered releases
- E Series router (ERX7xx models, ERX14xx models, the ERX310 router, the E120 router, or the E320 router)
- ASIC-based line modules that support Fast Ethernet or Gigabit Ethernet

Overview

Figure 23 on page 639 shows how you can use static subscriber interfaces to differentiate traffic for VPN access, based on the traffic's source address.

Figure 23: Subscriber Interfaces Using a Source Address to Demultiplex Traffic



Using a Source Address to Demultiplex Traffic

To configure the static subscriber interfaces shown in Figure 23 on page 639, perform the following steps:

1. Configure a primary IP interface on a supported layer 2 interface.
 - a. Create a layer 2 interface.


```
host1(config)#interface fastEthernet 4/1
```
 - b. Create a primary IP interface.


```
host1(config-if)#ip address 10.1.1.1 255.255.255.0
```
 - c. Exit Interface Configuration mode.


```
host1(config-if)#exit
```
2. Configure subscriber interface IP1.

- a. Create the shared IP interface.

```
host1(config)#virtual-router vra
  Proceed with new virtual-router creation? [confirm] yes
host1:vra(config)#interface ip ip1
```

Use the specified name to refer to the shared IP interface; you cannot use the layer 2 interface to refer to the shared IP interface, because the shared interface can be moved.

Use the **no** version to delete the IP interface.

- b. Associate the shared IP interface with the layer 2 interface by using one of the following methods:

- Static

```
host1:vra(config-if)#ip share-interface fastEthernet 4/1
```

If you issue this command on a shared IP interface, you cannot issue the **ip share-nexthop** command for the interface.

After creating the shared IP interface, you can configure it as you do any other IP interface. The shared interface is operationally up when the layer 2 interface is operationally up and IP is properly configured. You can create operational shared IP interfaces in the absence of a primary IP interface.

Use the **no** version to remove the association between the layer 2 interface and the shared IP interface. You can delete shared and primary IP interfaces independently.

- Dynamic

```
host1:vra(config-if)#ip share-nexthop 10.1.1.2
```

You can use this command to specify that the shared IP interface dynamically tracks a next hop. If the next hop changes, the shared IP interface moves to the new layer 2 interface associated with the IP interface toward the new next hop.

If you issue this command on a shared IP interface, you cannot issue the **ip share-interface** command for the interface. If you issue this command on a shared IP interface, the shared interface cannot dynamically track the next hop for the specified destination if the next-hop IP address is resolvable over MPLS. If you specify a virtual router, the command fails if the VR does not already exist. If you do not specify a VR, the current VR is assumed.

After creating the shared IP interface, you can configure it as you do any other IP interface. The shared interface is operationally up when the layer 2 interface associated with the specified next hop is operationally up and IP is properly configured.

Use the **no** version to halt tracking of the next hop.

- c. To fully configure the shared interface, assign an address or make it unnumbered.

```
host1:vra(config-if)#ip unnumbered loopback 0
```

- d. Specify the source addresses for the subscriber interface to use to demultiplex traffic, then exit Interface Configuration mode.


```
host1:vra(config-if)#ip source-prefix 10.10.3.0 255.255.255.0
host1:vra(config-if)#exit
```

On the ERX1440 router or the E320 router, you can configure up to 1024 subnets for static subscriber interfaces per primary IP interface when each subnet has a variable network mask that is less than /32. The number of subnets identifying a single route (/32) is still limited by the global maximum of 16,000 hosts per line module.

Use the **no** version to remove the association between the interface and the specified IP source address and mask.

3. Create a static route that sends traffic for destination address 10.10.3.0 to subscriber interface IP1.

```
host1:vra(config)#ip route 10.10.3.0 255.255.255.0 ip ip1
```

4. Repeat Step 2 to configure subscriber interface IP2.

```
host1(config)#virtual-router vrb
Proceed with new virtual-router creation? [confirm] yes
host1:vr(config)#interface ip ip2
host1:vr(config-if)#ip share-interface fastEthernet 4/1
host1:vr(config-if)#ip unnumbered loopback 0
host1:vr(config-if)#ip source-prefix 10.10.4.0 255.255.255.0
host1:vr(config-if)#exit
```

5. Create a static route that sends traffic for destination address 10.10.4.0 to subscriber interface IP2.

```
host1:vr(config)#ip route 10.10.4.0 255.255.255.0 ip ip2
```

6. Specify that DHCP relay does not install host routes—this avoids a conflict that can causes undesirable ARP behavior.

```
host1(config)#set dhcp relay inhibit-access-route-creation
```

For details about the cause of this conflict and the use of the **set dhcp relay inhibit-access-route-creation** command to avoid the conflict, see [“Configuring DHCP Relay Proxy” on page 540](#).

Related Documentation

- [Subscriber Interfaces Overview on page 619](#)
- [Characteristics of Subscriber Interfaces on page 623](#)
- [How MAC Address Validation State Inheritance Works on page 633](#)
- [Example: Configuring Dynamic Subscriber Interfaces on page 642](#)
- [Example: Configuring Dynamic Subscriber Interfaces Using Loopback Interfaces on page 650](#)
- *interface fastEthernet*
- *interface ip*
- *ip address*
- *ip demux-type da-prefix*

- *ip destination-prefix*
- *ip share-interface*
- *ip share-nexthop*
- *ip source-prefix*
- *ip unnumbered*
- *virtual-router*

Example: Configuring Dynamic Subscriber Interfaces

You can configure dynamic subscriber interfaces in the following configurations:

- IP over Ethernet
- IP over VLAN over Ethernet
- IP over bridged Ethernet over ATM
- GRE tunnels

The following topics describe how to create each of these basic configurations. In addition, [“Example: Configuring Dynamic Subscriber Interfaces Using Loopback Interfaces” on page 650](#), provides a detailed sample configuration.

- [Example: Configuring Dynamic Subscriber Interfaces over Ethernet on page 642](#)
- [Example: Configuring Dynamic Subscriber Interfaces over VLANs on page 644](#)
- [Example: Configuring Dynamic Subscriber Interfaces over Bridged Ethernet on page 646](#)
- [Example: Configuring Dynamic Subscriber Interfaces over GRE Tunnels on page 648](#)

Example: Configuring Dynamic Subscriber Interfaces over Ethernet

This example shows how you can configure a dynamic subscriber interface in an IP over Ethernet configuration by using DHCP events.

- [Requirements on page 642](#)
- [Overview on page 643](#)
- [Configuring Dynamic Subscriber Interfaces over Ethernet on page 643](#)

Requirements

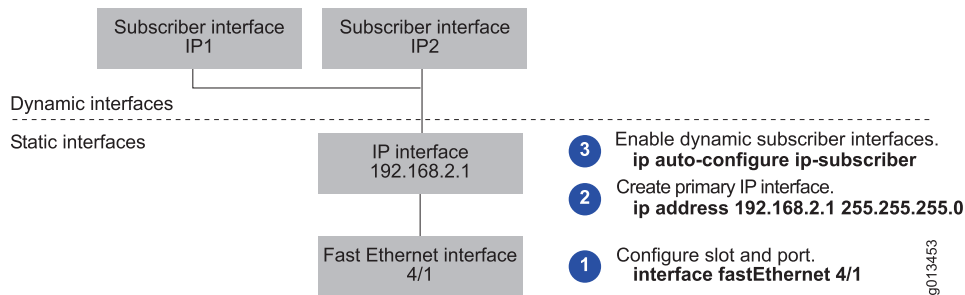
This example uses the following software and hardware components:

- JunosE Release 7.1.0 or higher-numbered releases
- E Series router (ERX7xx models, ERX14xx models, the ERX310 router, the E120 router, or the E320 router)
- ASIC-based line modules that support Fast Ethernet or Gigabit Ethernet

Overview

Figure 24 on page 643 shows the interface stack built for this configuration.

Figure 24: IP over Ethernet Dynamic Subscriber Interface Configuration



Configuring Dynamic Subscriber Interfaces over Ethernet

To configure a dynamic subscriber interface in an IP over Ethernet configuration by using DHCP events, perform the following steps:

1. Configure the DHCP server.

For instructions, see “Configuring the DHCP Local Server” on page 485.

2. Specify a Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet port.

```
host1(config)#interface fastEthernet 4/1
```

Use the **no** version to remove IP from an interface or subinterface. You must issue the **no** version from the highest level down; you cannot remove an interface or a subinterface if the one above it still exists.

3. Create the primary IP interface by assigning an IP address and mask to the Ethernet interface (or make it unnumbered).

```
host1(config-if)#ip address 192.168.2.1 255.255.255.0
```

You must specify the layer 2 encapsulation before you set the IP address. Issuing this command creates the primary IP interface. You must create a primary IP interface on which to enable dynamic creation of subscriber interfaces.

Use the **no** version to remove the IP address or to disable IP processing.

4. Configure the primary IP interface to enable dynamic creation of subscriber interfaces.

```
host1(config-if)#ip auto-configure ip-subscriber
```

The router creates the required dynamic subscriber interfaces when the IP address is assigned to the associated subscriber. The address might be assigned by an external DHCP server, the DHCP local server, or the packet detect feature.

Use the **include-primary** keyword to specify that the primary interface can be assigned to a subscriber. Use the **exclude-primary** keyword to specify that the primary interface is not used for subscribers. The primary interface is not assigned to a subscriber by default. You can issue this command from Interface Configuration mode, Subinterface Configuration mode, or Profile Configuration mode.

Use the **no** version to disable creation of dynamic subscriber interfaces associated with this primary IP interface. Use the **no** version with the **include-primary** keyword to specify that the primary interface is not assigned to a subscriber.

5. (Optional) Append the virtual router name to the subscriber interface in case of DSI configuration.

host1(config-if)#ip auto-configure append-virtual-router-name

You can use this command to allow more than one subscriber to have the same IP address across different virtual routers in the DSI configuration by appending the virtual router name to the interface. You can issue this command from either Interface Configuration mode or Profile Configuration mode.

Use the **no** version to disable ip auto-configure on the static primary interface if it is already configured. This feature is enabled by default in a non-DSI configuration with the DHCP local server.

6. (Optional) Specify the source address of traffic that is destined for the primary IP interface.

host1(config-if)#ip source-prefix 192.168.2.1 255.255.255.0

You can use this command to configure a subscriber interface or a primary IP interface enabled for dynamic creation of subscriber interfaces to demultiplex traffic with the specified source address. Use the **no** version to remove the association between the interface and the specified IP source address and mask.

Example: Configuring Dynamic Subscriber Interfaces over VLANs

This example shows how you can configure a dynamic subscriber interface in an IP over VLAN over Ethernet configuration by using DHCP events.

- [Requirements on page 644](#)
- [Overview on page 644](#)
- [Configuring Dynamic Subscriber Interfaces over VLANs on page 645](#)

Requirements


This example uses the following software and hardware components:

- JunosE Release 7.1.0 or higher-numbered releases
- E Series router (ERX7xx models, ERX14xx models, the ERX310 router, the E120 router, or the E320 router)
- ASIC-based line modules that support Fast Ethernet or Gigabit Ethernet

Overview

[Figure 25 on page 645](#) shows the interface stack built for this configuration.

Figure 25: IP over VLAN over Ethernet Dynamic Subscriber Interface Configuration



Configuring Dynamic Subscriber Interfaces over VLANs

To configure a dynamic subscriber interface in an IP over VLAN over Ethernet configuration by using DHCP events, perform the following steps:

1. Configure the DHCP server.

For instructions, see “[Configuring the DHCP Local Server](#)” on page 485.

2. Specify a Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet port.

```
host1(config)#interface gigabitEthernet 1/0
```



NOTE: You can configure only the primary port, 0, on the Gigabit Ethernet module. The router automatically uses the redundant port if the primary port fails.

Use the **no** version to remove IP from an interface. You must issue the **no** version from the highest level down; you cannot remove an interface or a subinterface if the one above it still exists.

3. Specify VLAN as the encapsulation method on the interface. This command creates the VLAN major interface.

```
host1(config-if)#encapsulation vlan
```

Use the **no** version to disable VLAN encapsulation on the interface.

4. Create a VLAN subinterface by adding a subinterface number to the interface identification command.

```
host1(config-if)#interface gigabitEthernet 1/0.1
```

5. Assign a unique VLAN ID to the VLAN subinterface.

```
host1(config-if)#vlan id 101
```

You need to specify a VLAN ID number that is in the range 0–4095 and is unique within the Ethernet interface. Issue the **vlan id** command before you configure any upper-layer interfaces, such as IP.

There is no **no** version.

6. Create the primary IP interface by assigning an IP address and mask to the VLAN subinterface (or make it unnumbered).

```
host1(config-if)#ip address 192.168.2.10 255.255.255.0
```

You must specify the layer 2 encapsulation before you set the IP address. Issuing this command creates the primary IP interface. You must create a primary IP interface on which to enable dynamic creation of subscriber interfaces.

Use the **no** version to remove the IP address or to disable IP processing.

7. (Optional) Append the virtual router name to the subscriber interface in case of DSI configuration.

```
host1(config-if)#ip auto-configure append-virtual-router-name
```

8. (Optional) Specify the source address of traffic that is destined for the primary IP interface.

```
host1(config-if)#ip source-prefix 192.168.2.10 255.255.255.0
```

You can use this command to configure a subscriber interface or a primary IP interface enabled for dynamic creation of subscriber interfaces to demultiplex traffic with the specified source address. Use the **no** version to remove the association between the interface and the specified IP source address and mask.

Example: Configuring Dynamic Subscriber Interfaces over Bridged Ethernet

This example shows how you can configure a dynamic subscriber interface in an IP over bridged Ethernet over ATM configuration by using DHCP events.

- [Requirements on page 646](#)
- [Overview on page 646](#)
- [Configuring Dynamic Subscriber Interfaces over Bridged Ethernet on page 647](#)

Requirements

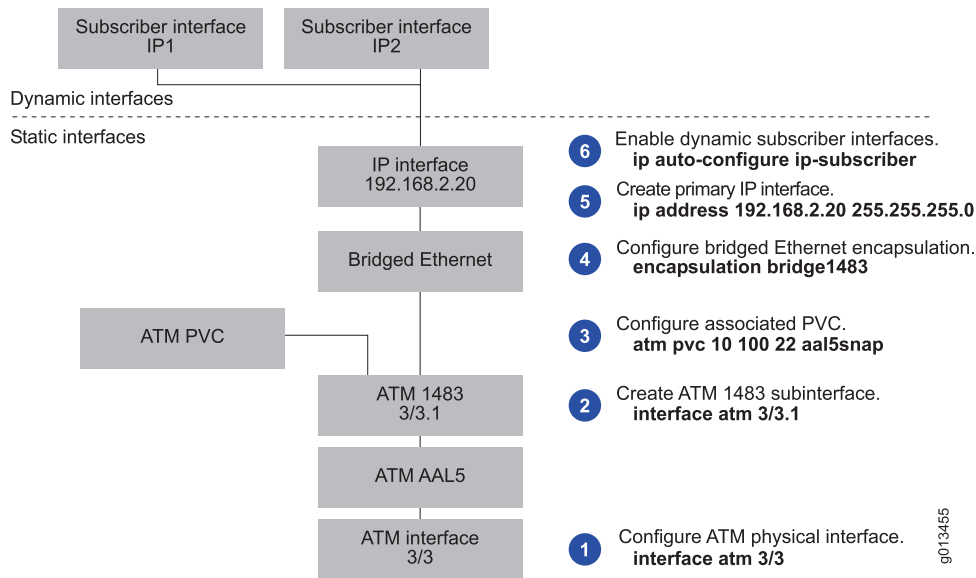
This example uses the following software and hardware components:

- JunosE Release 7.1.0 or higher-numbered releases
- E Series router (ERX7xx models, ERX14xx models, the ERX310 router, the E120 router, or the E320 router)
- ASIC-based line modules that support Fast Ethernet or Gigabit Ethernet

Overview

[Figure 26 on page 647](#) shows the interface stack built for this configuration.

Figure 26: IP over Bridged Ethernet over ATM Dynamic Subscriber Interface Configuration



Configuring Dynamic Subscriber Interfaces over Bridged Ethernet

To configure a dynamic subscriber interface in an IP over bridged Ethernet over ATM configuration by using DHCP events, perform the following steps:

1. Configure the DHCP server.
For instructions, see “Configuring the DHCP Local Server” on page 485.
2. Create an ATM major interface.
host1(config)#interface atm 3/3
3. Create an ATM 1483 subinterface.
host1(config-if)#interface atm 3/3.1
4. Configure an associated PVC for the ATM 1483 subinterface by specifying the VCD, the VPI, the VCI, and the encapsulation type.
host1(config-subif)#atm pvc 10 100 22 aal5snap
Specify the VCD, the VPI, the VCI, and the encapsulation type. (For more information about these parameters, see the *Creating a Basic Configuration* section in the *JunosE Link Layer Configuration Guide*.)
Use the **no** version to remove the specified PVC.
5. Specify bridged Ethernet as the encapsulation method on the ATM 1483 subinterface.
host1(config-subif)#encapsulation bridge1483
Use the **no** version to remove bridged Ethernet as the encapsulation method on the interface.
6. Create the primary IP interface by assigning an IP address and mask to the bridged Ethernet interface (or make it unnumbered).

```
host1(config-subif)#ip address 192.168.2.20 255.255.255.0
```

You must specify the layer 2 encapsulation before you set the IP address. Issuing this command creates the primary IP interface. You must create a primary IP interface on which to enable dynamic creation of subscriber interfaces.

Use the **no** version to remove the IP address or to disable IP processing.

7. Configure the primary IP interface to enable dynamic creation of subscriber interfaces.

```
host1(config-subif)#ip auto-configure ip-subscriber
```

8. (Optional) Append the virtual router name to the subscriber interface in case of DSI configuration.

```
host1(config-if)#ip auto-configure append-virtual-router-name
```

9. (Optional) Specify the source address of traffic that is destined for the primary IP interface.

```
host1(config-subif)#ip source-prefix 192.168.2.20 255.255.255.0
```

You can use this command to configure a subscriber interface or a primary IP interface enabled for dynamic creation of subscriber interfaces to demultiplex traffic with the specified source address. Use the **no** version to remove the association between the interface and the specified IP source address and mask.

Example: Configuring Dynamic Subscriber Interfaces over GRE Tunnels

This example shows how you can configure a dynamic subscriber interface in a GRE tunnel configuration by using packet detection.

- [Requirements on page 648](#)
- [Overview on page 648](#)
- [Configuring Dynamic Subscriber Interfaces over GRE Tunnels on page 649](#)

Requirements


This example uses the following software and hardware components:

- JunosE Release 7.1.0 or higher-numbered releases
- E Series router (ERX7xx models, ERX14xx models, the ERX310 router, the E120 router, or the E320 router)
- ASIC-based line modules that support Fast Ethernet or Gigabit Ethernet

Overview

[Figure 27 on page 649](#) shows the interface stack built for this configuration.

Figure 27: GRE Tunnel Dynamic Subscriber Interface Configuration



Configuring Dynamic Subscriber Interfaces over GRE Tunnels

To configure a dynamic subscriber interface in a GRE tunnel configuration by using packet detection, perform the following steps:

1. Create a GRE tunnel interface.

For instructions, see the *Configuration Tasks* section in the *JunosE IP Services Configuration Guide*.

2. Create the primary IP interface by assigning an IP address and mask to the bridged Ethernet interface (or make it unnumbered).

```
host1(config-subif)#ip address 192.168.2.20 255.255.255.0
```

You must specify the layer 2 encapsulation before you set the IP address. Issuing this command creates the primary IP interface. You must create a primary IP interface on which to enable dynamic creation of subscriber interfaces.

Use the **no** version to remove the IP address or to disable IP processing.

3. Configure the packet detect feature and specify that IP automatically detect packets that do not match any entries in the demultiplexer table.

```
host1(config-if)#ip auto-detect ip-subscriber
```

When an unmatched packet is detected, an event is generated that determines whether to create a dynamic subscriber interface.

Use the **no** version to restore the default, in which packet detection is disabled.

4. Configure the primary IP interface to enable dynamic creation of subscriber interfaces.

```
host1(config-subif)#ip auto-configure ip-subscriber
```

5. (Optional) Append the virtual router name to the subscriber interface in case of DSI configuration.

```
host1(config-if)#ip auto-configure append-virtual-router-name
```

6. (Optional) Specify the IP inactivity timer.

```
host1(config-subif)#ip inactivity-timer 100
```

A dynamically created subscriber interface is deleted if it is inactive for a period longer than the inactivity timer value. The timer value can be in the range 1–65335 minutes. A timer value of 0 specifies that dynamically created subscriber interfaces are never deleted by the inactivity timer.

Use the **no** version to restore the default, in which inactivity timer feature is disabled.

7. (Optional) Specify the source address of traffic that is destined for the primary IP interface.

```
host1(config-subif)#ip source-prefix 192.168.2.1 255.255.255.0
```

You can use this command to configure a subscriber interface or a primary IP interface enabled for dynamic creation of subscriber interfaces to demultiplex traffic with the specified source address. Use the **no** version to remove the association between the interface and the specified IP source address and mask.

Related Documentation

- [Dynamic Creation of Subscriber Interfaces Overview on page 628](#)
- [Example: Configuring Dynamic Subscriber Interfaces Using Loopback Interfaces on page 650](#)

Example: Configuring Dynamic Subscriber Interfaces Using Loopback Interfaces

This example shows how to configure dynamic subscriber interfaces by using the same loopback interface referenced by multiple unnumbered IP interfaces.

- [Requirements on page 650](#)
- [Overview on page 650](#)
- [Configuring a Dynamic Subscriber Interface on page 651](#)

Requirements

This example uses the following software and hardware components:

- JunosE Release 7.1.0 or higher-numbered releases
- E Series router (ERX7xx models, ERX14xx models, the ERX310 router, the E120 router, or the E320 router)
- ASIC-based line modules that support Fast Ethernet or Gigabit Ethernet

Overview

The procedure in this topic shows how to configure dynamic subscriber interfaces by using the same loopback interface referenced by multiple unnumbered IP interfaces. Instead of assigning a different IP address to each physical interface, this example assigns an IP address to a loopback interface (loopback 0). Each physical interface is then configured as an unnumbered IP interface, referencing the same loopback interface. This example uses a DHCP local server.

This approach has the following benefits:

- A loopback interface provides a stable IP address that can minimize the impact if a physical interface in the network goes down.
- Unnumbered IP interfaces preserve valuable IP address space.

Configuring a Dynamic Subscriber Interface

Step-by-Step Procedure

To configure dynamic subscriber interfaces, perform the following steps:

1. Enable the DHCP local server for standalone mode.

```
host1(config)#service dhcp-local standalone
```

2. Configure DHCP relay to use information in the giaddr in DHCP server-destined packets to identify the primary interface on which dynamic subscriber interfaces are built. See the *Using the Giaddr to Identify the Primary Interface for Dynamic Subscriber Interfaces* section in “DHCP Relay and BOOTP Relay Overview” on page 512 for additional information about this feature.

```
host1(config)#set dhcp relay giaddr-selects-interface
```

Use the **no** version to restore the default in which DHCP relay builds dynamic subscriber interfaces on the IP interface that is used for DHCP server-destined messages.

3. Access DHCP Local Pool Configuration mode for the local address pool.

```
host1(config)#ip dhcp-local pool ispWestford
```

The DHCP local server uses pool names other than default to maintain configuration information for subscribers to a particular domain.

Use the **no** version to prevent the DHCP local server from supplying IP addresses from the specified pool.

4. Specify the enduring IP addresses that the DHCP local server can assign from the local address pool.

```
host1(config-dhcp-local)#network 10.20.0.0 255.255.192.0
```

5. Specify the router to forward traffic from the IP addresses to destinations on other subnets.

```
host1(config-dhcp-local)#default-router 10.20.32.1
```

You can specify the IP address of a primary server, and optionally, specify the IP address of a secondary server.

Use the **no** version to remove the association between the address pool and the router.

6. Exit DHCP Local Pool Configuration mode.

```
host1(config-dhcp-local)#exit
```

7. Configure a loopback interface.

```
host1(config)#interface loopback 0
```



BEST PRACTICE: We recommend that you configure a 32-bit subnet mask for the loopback interface. For example, if you configure a loopback interface with the IP address and mask as 1.1.1.1/16, the 1.1.0.0/16 route entry is entered on the line module and all traffic destined to the 1.1.0.0/16 subnet is forwarded to the SRP module by the line module. Although the SRP module responds only to traffic destined to the 1.1.1.1 subnet and discards traffic to all other host IP addresses within that subnet (1.1.1.1/16), if no specific or longer route entry is found or if the SRP module receives too much traffic from subnets other than 1.1.1.1, the CPU utilization on the SRP module reaches the saturation level.

If you use a subnet mask other than a /32 mask for the IP address configured on the loopback interface, traffic from the entire subnet is routed to the loopback interface. Therefore, that subnet cannot be routed through any other interface on the router, unless a more specific route points to another interface.

8. Assign an IP address and mask to the loopback interface.

```
host1(config-if)#ip address 10.20.32.1 255.255.255.255
```

You must specify the layer 2 encapsulation before you set the IP address. Issuing this command creates the primary IP interface. You must create a primary IP interface on which to enable dynamic creation of subscriber interfaces.

Use the **no** version to remove the IP address or to disable IP processing.

9. Exit Interface Configuration mode.

```
host1(config-if)#exit
```

10. Specify a Fast Ethernet port.

```
host1(config)#interface fastEthernet 3/0
```

Use the **no** version to remove IP from an interface or subinterface. You must issue the **no** version from the highest level down; you cannot remove an interface or a subinterface if the one above it still exists.

11. Create an unnumbered primary IP interface associated with the loopback interface configured in Steps 6 and 7.

```
host1(config-if)#ip unnumbered loopback 0
```

This command enables IP processing on an interface without assigning an explicit IP address to the interface. You must specify an interface location, which is the identifier of another interface on which the router has an assigned IP address. This interface cannot be another unnumbered interface.

Use the **no** version to disable IP processing on the interface.

12. Configure a static primary IP interface to use framed routes as source IP addresses when creating dynamic subscriber interfaces.

```
host1(config-if)#ip use-framed-routes ip-subscriber
```

The router uses the Framed-Route RADIUS attribute [22] sent in Access-Accept messages to apply framed routes to subscriber interfaces associated with the primary interface.

Use the **no** version to disable the use of framed routes when creating dynamic subscriber interfaces associated with this primary IP interface.

13. Configure the primary IP interface to enable dynamic creation of subscriber interfaces. The specified IP interface is considered the primary interface.

```
host1(config-if)#ip auto-configure ip-subscriber
```

The router creates the required dynamic subscriber interfaces when the IP address is assigned to the associated subscriber. The address might be assigned by an external DHCP server, the DHCP local server, or the packet detect feature.

Use the **include-primary** keyword to specify that the primary interface can be assigned to a subscriber. Use the **exclude-primary** keyword to specify that the primary interface is not used for subscribers. The primary interface is not assigned to a subscriber by default. You can issue this command from Interface Configuration mode, Subinterface Configuration mode, or Profile Configuration mode.

Use the **no** version to disable creation of dynamic subscriber interfaces associated with this primary IP interface. Use the **no** version with the **include-primary** keyword to specify that the primary interface is not assigned to a subscriber.

14. (Optional) Append the virtual router name to the subscriber interface in case of DSI configuration.

```
host1(config-if)#ip auto-configure append-virtual-router-name
```

You can use this command to allow more than one subscriber to have the same IP address across different virtual routers in the DSI configuration by appending the virtual router name to the interface. You can issue this command from either Interface Configuration mode or Profile Configuration mode.

Use the **no** version to disable ip auto-configure on the static primary interface if it is already configured. This feature is enabled by default in a non-DSI configuration with the DHCP local server.

15. Exit Interface Configuration mode.

```
host1(config-if)#exit
```

16. Repeat Steps 9 through 12 for each Fast Ethernet interface on which you want to configure dynamic subscriber interfaces. For example:

```
host1(config)#interface fastEthernet 3/1
host1(config-if)#ip unnumbered loopback 0
host1(config-if)#ip auto-configure ip-subscriber
host1(config-if)#ip auto-configure append-virtual-router-name
host1(config-if)#exit
host1(config)#interface fastEthernet 3/2
host1(config-if)#ip unnumbered loopback 0
host1(config-if)#ip auto-configure ip-subscriber
host1(config-if)#ip auto-configure append-virtual-router-name
host1(config-if)#exit
```

**Related
Documentation**

- [Example: Configuring Dynamic Subscriber Interfaces on page 642](#)
- *atm pvc*
- *encapsulation bridge1483*
- *encapsulation vlan*
- *interface atm*
- *interface fastEthernet*
- *interface gigabitEthernet*
- *interface tenGigabitEthernet*
- *interface loopback*
- *ip address*
- *ip auto-configure append-virtual-router-name*
- *ip auto-configure ip-subscriber*
- *ip auto-detect ip-subscriber*
- *ip dhcp-local pool*
- *ip inactivity-timer*
- *ip source-prefix*
- *ip unnumbered*
- *ip use-framed-routes ip-subscriber*
- *network*
- *service dhcp-local*
- *set dhcp relay giaddr-selects-interface*
- *vlan id*

Monitoring Subscriber Interfaces

This chapter describes how to monitor static and dynamic subscriber interfaces for remote access to the E Series router. This chapter contains the following sections:

- [Monitoring Subscriber Interfaces Overview on page 655](#)
- [Monitoring Subscriber Interfaces on page 655](#)
- [Monitoring Active IP Subscribers Created by Subscriber Management on page 656](#)

Monitoring Subscriber Interfaces Overview

- Purpose** The state of the subscriber interface is determined by state of the Ethernet interface and the existence of the primary IP interface, which you can monitor with the *show ip interface* command. For information about using the **show ip interface** command, see the *Monitoring IP* section in the *JunosE IP Services Configuration Guide*.
- Action** You can use the **show ip demux interface** command to monitor the configuration of subscriber interfaces.

Monitoring Subscriber Interfaces

- Purpose** Display information about subscriber interfaces.

- Action** To display subscriber interface information:

```
host1#show ip demux interface fastEthernet 2/0
Prefix/Length  SA/DA  Subscriber-Intf  VR/VRF  Description
10.12.2.2/32   SA     ip subsc1       3        subsc1@test
10.12.2.3/32   SA     ip subsc2       3        subsc2@test
10.12.2.4/32   SA     ip subsc3       3        subsc3@test
10.12.2.5/32   SA     ip subsc4       3        subsc4@test
```

- Meaning** [Table 147 on page 655](#) lists the **show ip demux interface** command output fields.

Table 147: show ip demux interface Output Fields

Field Name	Field Description
Prefix/Length	Source or destination addresses that the subscriber interface demultiplexes
SA/DA	Demultiplexing method for subscriber interface

Table 147: show ip demux interface Output Fields (*continued*)

Field Name	Field Description
SA	Source address
DA	Destination address
Subscriber-Intf	Name of shared interface on which subscriber interface is configured
VR/VRF	Name of virtual router (VR) or VPN routing and forwarding (VRF) instance on which the subscriber interface is configured
Description	Text description for the IP interface on which subscriber interface is configured (added with the <i>ip description</i> command)

Related Documentation

- *show ip demux interface*

Monitoring Active IP Subscribers Created by Subscriber Management

Purpose Display information about active IP subscribers that were created by the JunosE Software's subscriber management feature.

Action To display information about subscribers that were created by subscriber management:

```
host1# show ip-subscriber 2835349506
```

Id	User Name	Ip Address	Virtual Router	Interface
2835349506	user1@isp1.com	192.168.0.1	default	ip192.168.0.1
Id	Login time			
2835349506	WED AUG 23 20:46:24 2006			

```
host1# show ip-subscriber detail
```

Subscriber List				
Id	User Name	Ip Address	Virtual Router	Interface
2835349506	user1@isp1.com	192.168.0.1	default	ip192.168.0.1
Id	Login Time	Mac Address		Profile Handle
2835349506	WED AUG 23 20:46:24 2006	3000.0001.9365		13631489
Id	Interface Profile	Service Profile	Option 82	
2835349506	myProfile	profile22	FastEthernet 3/1	

Meaning [Table 148 on page 657](#) lists the **show ip-subscriber** command output fields.

Table 148: show ip-subscriber Output Fields

Field Name	Field Description
Id	ID of the subscriber
Username	Username used to retrieve information from RADIUS for the subscriber interface
Ip Address	IP address of the subscriber interface
Virtual Router	Name of the virtual router on which the subscriber interface is configured
Interface	Name of subscriber interface; ip indicates that subscriber manager created this interface
Login Time	Day, date, and time that the subscriber logged in
Mac Address	MAC address of the subscriber
Profile Handle	AAA profile handle
Interface Profile	Interface profile name used to configure the subscriber interface
Service Profile	IP service profile name used by subscriber management to authorize and configure the subscriber interface with AAA
Option 82	DHCP relay agent information (option 82) circuit identifier that describes the physical interface location associated with the subscriber

Related Documentation

- *show ip-subscriber*

PART 6

Managing Subscriber Services

- [Configuring Service Manager on page 661](#)
- [Monitoring Service Manager on page 749](#)

CHAPTER 28

Configuring Service Manager

This chapter describes how to use the Service Manager application to define, activate, and monitor networking services for your subscribers. This chapter discusses the following topics:

- [Service Manager Overview on page 663](#)
- [Service Manager Terms and Acronyms on page 663](#)
- [Service Manager Platform Considerations on page 664](#)
- [Service Manager References on page 665](#)
- [Configuring Service Manager on page 665](#)
- [Service Definitions Overview on page 667](#)
- [Creating Service Definitions on page 668](#)
- [Managing Your Service Definitions on page 671](#)
- [Referencing Policies in Service Definitions Overview on page 673](#)
- [Referencing QoS Configurations in Service Definitions on page 673](#)
- [Specifying QoS Profiles in Service Definitions on page 674](#)
- [Specifying QoS Parameter Instances in a Service Definition on page 675](#)
- [Modifying QoS Configurations with Service Manager on page 677](#)
- [Removing QoS Configurations Referenced by Service Manager on page 681](#)
- [QoS for Service Manager Considerations on page 682](#)
- [Configuring the Service Manager License on page 683](#)
- [Managing and Activating Service Sessions Overview on page 684](#)
- [Managing Subscriber Service Sessions by Using RADIUS Overview on page 684](#)
- [Activating Subscriber Service Sessions Using RADIUS on page 685](#)
- [Understanding Service Manager RADIUS Attributes on page 686](#)
- [Deactivating Service Sessions Using RADIUS on page 690](#)
- [Activating and Deactivating Subscriber Services by Using Mutex Groups Overview on page 691](#)
- [Activating and Deactivating Multiple Services on page 692](#)
- [Configuring a Mutex Service on page 693](#)

- [Combined and Independent IPv4 and IPv6 Services in a Dual Stack Overview on page 694](#)
- [Activation and Deactivation of IPv4 and IPv6 Services in a Dual Stack on page 696](#)
- [Understanding RADIUS Accounting for Service Manager on page 697](#)
- [Service Interim Accounting Overview on page 699](#)
- [Configuring Service Interim Accounting on page 702](#)
- [Configuring Calculation of Service Session Accounting Based on Scheduler Profiles Instead of Rate-Limit Profiles in Hierarchical Parent Groups for Forwarded Packets on page 703](#)
- [Service Interim Accounting for IPv4 and IPv6 Services in a Dual Stack Overview on page 705](#)
- [Collecting IPv6 Service Accounting Statistics Separately for Dual-Stack Subscribers Overview on page 706](#)
- [Managing Subscriber Service Sessions by Using the CLI Overview on page 707](#)
- [Activating Subscriber Service Sessions by Using the CLI Overview on page 708](#)
- [Activating Subscriber Sessions by Using the CLI on page 709](#)
- [Preprovisioning Service Sessions on page 711](#)
- [Service Session Profiles Overview on page 711](#)
- [Working with Service Session Profiles on page 712](#)
- [Deactivating Subscriber Service Sessions by Using the CLI Overview on page 715](#)
- [Gracefully Deactivating Subscriber Service Sessions on page 716](#)
- [Forcing Immediate Deactivation of Subscriber Service Sessions on page 717](#)
- [Using Service Session Profiles to Deactivate Service Sessions on page 717](#)
- [Configuring Service Manager Statistics on page 718](#)
- [Service Manager Performance Considerations on page 723](#)
- [Example: Tiered Service Definition on page 724](#)
- [Example: Video-on-Demand Service Definition on page 726](#)
- [Example: Voice-over-IP Service Definition on page 728](#)
- [Guided Entrance Service Overview on page 730](#)
- [Example: Guided Entrance Service Definition on page 731](#)
- [Using COA Messages with Guided Entrance Services on page 734](#)
- [Configuring the HTTP Local Server to Support Guided Entrance on page 735](#)
- [Redirection of Subscriber Sessions When HTTP Local Server is Disabled or Not Configured on page 739](#)
- [Example: Combined IPv4 and IPv6 Service in a Dual Stack Service Definition on page 740](#)
- [Preservation of the Original URL During Redirection of Subscriber Sessions on page 746](#)
- [Configuring the Preservation of the Original URL During Redirection of Subscriber Sessions on page 747](#)

Service Manager Overview

The JunosE Service Manager application provides authentication, service selection, and service activation and deactivation to subscribers. The application also collects accounting information and statistics, and monitors subscriber and service sessions.

Service Manager supports two client types—RADIUS and CLI. Service Manager starts when it receives a request from a RADIUS or CLI client. For RADIUS clients, RADIUS Access-Accept messages and Change-of-Authorization-Request (COA-Request) messages can create and delete Service Manager subscriber sessions and activate and deactivate service sessions. For CLI clients, CLI commands create and delete the subscriber sessions and activate and deactivate service sessions.

A subscriber's service is based on a service definition—service definitions can include profiles, policies, and quality of service (QoS) settings that define the scope of a service granted to the subscriber. Service definitions can also specify statistics configurations.

Service Manager provides convenience and flexibility to both service providers and subscribers.

- Providers are able to separate services and access technology and also to eliminate unprofitable flat-rate billing. They gain the ability to efficiently design, manage, and deliver services that subscribers want, and then bill subscribers based on connect time, bandwidth, and the actual service used.
- Subscribers benefit by gaining access to multiple simultaneous services—subscribers can dynamically connect to and disconnect from the services, when they want and for how long they want. They are billed based on the service type and usage, rather than being charged a set rate regardless of usage.

Related Documentation

- [Service Manager Platform Considerations on page 664](#)
- [Service Definitions Overview on page 667](#)
- [Configuring the Service Manager License on page 683](#)
- [Configuring Service Manager on page 665](#)

Service Manager Terms and Acronyms

Table 149 on page 663 defines terms and acronyms that are used in this discussion of the Service Manager application.

Table 149: Service Manager Terms and Acronyms

Term	Definition
Guided entrance	A service that creates a controlled Internet browsing environment by transparently directing the subscriber to a specific website. At the website, the subscriber is presented with a selection of available services. Also called <i>walled gardens</i> or <i>captive portals</i> .

Table 149: Service Manager Terms and Acronyms (*continued*)

Term	Definition
Macro language	The JunosE macro language that you use for service definitions
Mutex service	A service session that is part of a mutex group—the service definition for the service includes the mutex-group attribute.
RADIUS login method	The method that uses RADIUS VSAs in the Access-Accept packet to create a subscriber session and activate a service session when the subscriber logs in
RADIUS COA method	The method that uses RADIUS COA-Request messages and VSAs to create a subscriber session and activate a service session for a subscriber that is already logged in
Service definition	A macro file that defines a named parameterized description of a service; used to create a service instance and the resulting subscriber service session; can include a combination of parameters such as policy lists, rate-limit profiles, QoS profiles, and interface profiles
Service instance	An instance that is created when you specify parameter values for a service definition to create a service session
Service session	A session that is created when a service instance is activated for a subscriber; a subscriber can have multiple active service sessions
Service session profile	A provider-configured profile that applies optional attributes to a service session; CLI only

Related Documentation

- [Service Manager Overview on page 663](#)
- [Service Definitions Overview on page 667](#)
- [Understanding Service Manager RADIUS Attributes on page 686](#)
- [Service Session Profiles Overview on page 711](#)
- [Managing Subscriber Service Sessions by Using RADIUS Overview on page 684](#)

Service Manager Platform Considerations

Service Manager is supported on all E Series routers. For information about the modules supported on E Series routers:

- See the *ERX Module Guide* for modules supported on ERX7xx models, ERX14xx models, and the ERX310 Broadband Services Router.
- See the *E120 and E320 Module Guide* for modules supported on the E120 and E320 Broadband Services Routers.

- Related Documentation**
- [Service Manager Overview on page 663](#)
 - [Service Manager References on page 665](#)
 - [Configuring the Service Manager License on page 683](#)

Service Manager References

For more information about the topics covered in this chapter, see the following documents:

- Data-Over-Cable Service Interface Specifications (DOCSIS) 2.0 Radio Frequency Interface Specification CM-SP-RFv2.0-I10-051209.
- For information about using the JunosE Software's macro language, see the *Writing CLI Macros* chapter in the *JunosE System Basics Configuration Guide*.
- For information about setting up policy-based routing features for Service Manager, such as rate-limit profiles, classifier control lists, policy lists, and hierarchical and merged policies, see the *JunosE Policy Management Configuration Guide*.
- For information about creating QoS profiles and QoS parameters, see the *JunosE Quality of Service Configuration Guide*.
- For information about creating IPv4 interface profiles, see the *Configuring IP* chapter in the *JunosE IP, IPv6, and IGP Configuration Guide*.

- Related Documentation**
- [Service Manager Overview on page 663](#)
 - [Service Manager Terms and Acronyms on page 663](#)
 - [Service Manager Platform Considerations on page 664](#)
 - [Service Definitions Overview on page 667](#)
 - [Configuring Service Manager on page 665](#)

Configuring Service Manager


To use the Service Manager application to create subscriber service sessions, you perform the following tasks:

- Create and manage service definitions. (See “[Managing Your Service Definitions](#)” on [page 671](#).)
 - Use the macro language to define service definitions.
 - Download service definition macro files to the router's nonvolatile storage (NVS).
 - Install service definitions on the router.
 - Uninstall service definitions.
- Configure the Service Manager license. (See “[Configuring the Service Manager License](#)” on [page 683](#).)

- Configure RADIUS accounting. (See [“Understanding RADIUS Accounting for Service Manager” on page 697](#) and [“Configuring Service Interim Accounting” on page 702.](#))
- Use RADIUS login and RADIUS COA to manage subscriber service sessions. (See [“Managing Subscriber Service Sessions by Using RADIUS Overview” on page 684](#) and [“Activating Subscriber Service Sessions Using RADIUS” on page 685.](#))
 - Specify the subscriber. (See [“Understanding Service Manager RADIUS Attributes” on page 686.](#))
 - Specify optional attributes.
 - Enable statistics collection. (See [“Configuring Service Manager Statistics” on page 718.](#))
 - Activate the service session. (See [“Deactivating Service Sessions Using RADIUS” on page 690.](#))
 - Deactivate service sessions. (See [“Activating Subscriber Service Sessions Using RADIUS” on page 685.](#))
 - (Optional for RADIUS COA method) Configure the COA feature for the RADIUS dynamic-request server. (See the *Configuring RADIUS Dynamic-Request Server* chapter in this guide)
- Use the CLI to manage subscriber service sessions. (See [“Preprovisioning Service Sessions” on page 711](#) and [“Activating Subscriber Sessions by Using the CLI” on page 709.](#))
 - Specify the subscriber. (See [“Working with Service Session Profiles” on page 712.](#))
 - Create and apply optional service session profiles.
 - Enable statistics collection. (See [“Configuring Service Manager Statistics” on page 718.](#))
 - Activate the service session. (See [“Activating Subscriber Sessions by Using the CLI” on page 709.](#))
 - Deactivate service sessions. (See [“Deactivating Subscriber Service Sessions by Using the CLI Overview” on page 715.](#))

[Figure 28 on page 667](#) shows the sequence of operations you use to create and monitor subscriber service sessions.

Figure 28: Service Manager Configuration Flowchart



- Related Documentation**
- [Service Manager Overview on page 663](#)
 - [Service Definitions Overview on page 667](#)
 - [Creating Service Definitions on page 668](#)

Service Definitions Overview

A service definition is a high-level, platform-independent template that defines a service that you want to let your subscribers use. You use the JunosE Software's embedded macro language on your computer to create the macro file that defines the service. You copy and install the macro file on the E Series Broadband Services Routers, and then you can associate the service definition with subscribers to create their service sessions.

Service definitions gives you flexibility by enabling you to use:

- A single service definition to create a service for multiple subscribers.
- Parameterized service definitions to create variations of a service definition.
- Different service definitions to create multiple services for a single subscriber.

A service definition might use the following types of JunosE objects to define the characteristics and capabilities of the service you want to provide:

- Interface profiles—Specify a set of characteristics that can be dynamically assigned to IP interfaces. A service definition must use at least one interface profile.
- Policy lists—Specify policy actions for traffic traversing an interface.
- Classifier lists—Specify the criteria by which the router defines a packet flow.
- Rate-limit profiles—Specify a set of bandwidth attributes and associated actions that limit a classified packet flow or a source interface to a rate that is less than the physical rate of the port.
- QoS parameters—Specify attributes such as shaping rate, shared-shaping rate, assured rate, and scheduler weight for scheduler nodes and queues.
- QoS profiles—Specify queue, drop statistics gathering, and scheduler configuration for an interface hierarchy.

**Related
Documentation**

- [Creating Service Definitions on page 668](#)
- [Managing Your Service Definitions on page 671](#)
- [Referencing Policies in Service Definitions Overview on page 673](#)
- [Referencing QoS Configurations in Service Definitions on page 673](#)
- [Specifying QoS Profiles in Service Definitions on page 674](#)

Creating Service Definitions

To create a service definition, you use the JunosE Software's macro language to specify the parameters that define the desired service. A macro file can define only one service—however, the file can have multiple templates to define characteristics of the desired service. You create service definitions independent of the Service Manager commands and operations, which are performed on the E Series router.

For detailed information about the JunosE Software's macro language, see the *Command Line Interface* chapter in the *JunosE System Basics Configuration Guide*.

[Figure 29 on page 670](#) is an example of a service definition macro file that creates a tiered service. A tiered service typically provides set bandwidths for both inbound and outbound traffic for a subscriber. In this example, the input (inputBW) and output (outputBW) bandwidth values are parameterized. This example assumes that QoS profile triplePlayIP and QoS parameter maxSubscBW are configured. See [“Example: Video-on-Demand Service Definition” on page 726](#), [“Example: Voice-over-IP Service Definition” on page 728](#), and [“Example: Guided Entrance Service Definition” on page 731](#) for additional service definition examples.

Service Manager only tracks JunosE objects that are passed back in the env.setResult method when a service definition is executed. [Table 150 on page 669](#) describes the supported objects:

Table 150: JunosE Objects Tracked by Service Manager

Name	Requirement	Description
input-stat-clacl	Optional	<ul style="list-style-type: none"> Collects input statistics from policy manager Can be a list of clacl
secondary-input-stat-clacl	Optional	<ul style="list-style-type: none"> Collects input statistics from policy manager Can be a list of clacl
output-stat-clacl	Optional	<ul style="list-style-type: none"> Collects output statistics from policy manager Can be a list of clacl
activate-profile	Required	<ul style="list-style-type: none"> Specifies the interface profile used on activation of the service Deletion of the profile is Service Manager's responsibility
deactivate-profile	Optional	<ul style="list-style-type: none"> Specifies the interface profile used on deactivation of the service If not specified, is the same as the activation-profile Deletion of the profile is Service Manager's responsibility
command-in-error	Optional	<ul style="list-style-type: none"> Passes the value env.getErrorCommand Service Manager displays the line in the service definition that has the error
command-error-status	Optional	<ul style="list-style-type: none"> Passes the value env.getErrorStatus Service Manager displays the error status for the error
service-interface-type	<ul style="list-style-type: none"> Optional for IPv4 or L2TP services Mandatory for independent IPv6 services or combined IPv4 and IPv6 services in a dual stack 	<ul style="list-style-type: none"> Specifies the type of interface, IPv4, IPv6, combination of IPv4 and IPv6, or L2TP, to which the service session profile must be applied
input-stat-epg	Optional	<ul style="list-style-type: none"> Collects input statistics associated with the external group from policy manager Both the external parent group and the corresponding hierarchical policy parameter must be specified Can be multiple pairs of external parent groups and hierarchical policy parameters

Table 150: JunosE Objects Tracked by Service Manager (*continued*)

Name	Requirement	Description
output-stat-epg	Optional	<ul style="list-style-type: none"> Collects output statistics associated with the external group from policy manager Both the external parent group and the corresponding hierarchical policy parameter must be specified Can be multiple pairs of external parent groups and hierarchical policy parameters
secondary-input-stat-epg	Optional	<ul style="list-style-type: none"> Collects input statistics associated with the external group that is attached at the secondary input stage from policy manager Both the external parent group and the corresponding hierarchical policy parameter must be specified Can be multiple pairs of external parent groups and hierarchical policy parameters

Figure 29: Sample Service Definition Macro File

```

!parameterizes input and output bandwidth
<# tiered(inputBW, outputBW) #>

<# uid := app.servicemanager.getUniqueId #>
<# name := "SM-tiered-" $ uid #>
<# oname := "SM-O-tiered-" $ uid #>

classifier-list matchAll ip any any
rate-limit-profile <# name #> one-rate
    committed-rate <# inputBW; '\n' #>

policy-list <# name; '\n' #>
    classifier-group matchAll precedence 10000
    rate-limit-profile <# name; '\n' #>
    traffic-class best-effort

policy-list <# oname; '\n' #>
    classifier-group matchAll precedence 10000
    traffic-class best-effort

profile <# name; '\n' #>
    ip policy secondary-input <# name #> statistics enabled merge
    ip policy output <# oname #> statistics enabled merge
    qos-profile triplePlayIP
    qos-parameter maxSubscBW <# outputBW; '\n' #>

<# env.setResult("activate-profile", name) #>
<# env.setResult("secondary-input-stat-clacl", "matchAll") #>
<# env.setResult("output-stat-clacl", "matchAll") #>

<# endtmp1 #>

```

Related Documentation

- [Service Definitions Overview on page 667](#)

- [Managing Your Service Definitions on page 671](#)
- [Referencing Policies in Service Definitions Overview on page 673](#)
- [Referencing QoS Configurations in Service Definitions on page 673](#)
- [Specifying QoS Profiles in Service Definitions on page 674](#)

Managing Your Service Definitions

After you have created the macro file for your service definition, you can perform the following operations with the service definition macro file:



NOTE: All new service sessions will be activated using the new service definition. Any existing service sessions that were activated using the original service definition continue to use the original definition until you deactivate the service session.

- [Copying a Service Definition Macro File on page 671](#)
- [Installing a Service Definition File on page 671](#)
- [Uninstalling a Service Definition File on page 672](#)
- [Updating an Existing Service Definition File on page 672](#)

Copying a Service Definition Macro File

You must copy the service definition from the local computer that you used to create the macro file to the router's NVS card.

To copy a service definition macro file from your computer to the router's NVS:

- Issue the **copy** command in Privileged Exec mode. You must specify the directory containing the macro file you want to copy and the name you want to use for the file in NVS.

```
host1#copy boston:/serviceDefs/triplePlay/tiered.mac tiered.mac
```

There is no **no** version.

Installing a Service Definition File

You must install the service definition before you can use it to create a service session. During installation, Service Manager precompiles the definition and extracts and extracts the definition file's timestamp. Precompiling the service definition improves Service Manager performance. The timestamp enables the Service Manager application to track any modifications you might make while the definition is being used.

To install a service definition file:

- Issue the **service-management install** command in Global Configuration mode.

```
host1(config)#service-management install tiered.mac
```

You must include the .mac extension while you use the **service-management install** command to install a service definition. After you install the service definition, you can use the definition to create service sessions for subscribers.

Use the **no** version of this command to remove a previously configured service definition.

Uninstalling a Service Definition File

You can uninstall a service definition file, for example, if you no longer want to use that definition. When you uninstall a service definition file, any existing service sessions that were activated using the original service definition continue to use the original definition until you deactivate the service session.

To uninstall a service definition file:

- Issue the **no service-management install** command in Global Configuration mode.

```
host1(config)#no service-management install tiered.mac
```

Updating an Existing Service Definition File

You can update an existing service definition file at any time. To update a service definition file:

To update an existing service definition, you make changes to the original macro file on your computer, copy the updated file to NVS, and install the updated file. All subsequent service sessions use the new service definition file. However, currently active service sessions continue to use the original definition file until the sessions are deactivated, then reactivated.

1. Use your text editor on your computer to make changes to the original service definition file.
2. Copy the updated service definition file back to your router's NVS—this overwrites the original file on the router.

```
! update the original macro file on the remote system
```

```
! copy the updated macro file to the router  
host1#copy boston:/serviceDefs/triplePlay/tiered.mac tiered.mac
```

3. Enter the Global Configuration mode and Install the new service definition file.

```
host1#configure terminal  
! install the updated service definition on the router  
host1(config)#service-management install tiered.mac
```

Related Documentation

- [Service Definitions Overview on page 667](#)
- [Creating Service Definitions on page 668](#)
- [Referencing QoS Configurations in Service Definitions on page 673](#)
- [Specifying QoS Profiles in Service Definitions on page 674](#)
- [Example: Tiered Service Definition on page 724](#)

- [Example: Voice-over-IP Service Definition on page 728](#)
- [Example: Video-on-Demand Service Definition on page 726](#)
- [Example: Guided Entrance Service Definition on page 731](#)
- *copy*
- *service-management install*

Referencing Policies in Service Definitions Overview

In Profile Configuration mode, policy interface commands for IP and L2TP allow attachments to be merged into any existing merge-capable attachment at an attachment point. Merged policies are dynamically created. Service Manager can request that multiple interface profiles be applied or removed at an interface as part of service activation or deactivation. Service Manager also specifies whether or not the attachments created from these interface profiles persist on subsequent reloads.

Service Manager can specify whether a component policy attachment is non-volatile. If the interface where the component policy is attached is volatile, then policy management makes the attachment volatile even when the Service Manager specifies otherwise. A non-volatile interface can have both volatile and non-volatile component policy attachments. The merged policy that is created is the merge of all component policies attached at a given attachment point regardless of their volatility. The merged policy and its attachments are always volatile and reconstructed on each reload operation.

For further details on merging policies, see the *Merging Policies* chapter in the *JunosE Policy Management Configuration Guide*.

Related Documentation

- [Creating Service Definitions on page 668](#)
- [Managing Your Service Definitions on page 671](#)
- [Referencing QoS Configurations in Service Definitions on page 673](#)
- [Modifying QoS Configurations with Service Manager on page 677](#)
- [Removing QoS Configurations Referenced by Service Manager on page 681](#)
- [Specifying QoS Profiles in Service Definitions on page 674](#)

Referencing QoS Configurations in Service Definitions

You can use QoS profiles and QoS parameters to define a service for a subscriber. For example, you can configure the shaping rate for traffic in a video service by using a QoS parameter instance.

To transmit the QoS configuration to the subscriber interface (that is, the forwarding interface at the top of the interface column), you must configure the QoS profiles and QoS parameter instances in static profiles.

To reference QoS configurations in service definitions, perform the following tasks:

**Related
Documentation**

- [Specifying QoS Profiles in Service Definitions on page 674](#)
- [Specifying QoS Parameter Instances in a Service Definition on page 675](#)
- [Creating Service Definitions on page 668](#)
- [Managing Your Service Definitions on page 671](#)
- [Referencing Policies in Service Definitions Overview on page 673](#)
- [Modifying QoS Configurations with Service Manager on page 677](#)
- [Removing QoS Configurations Referenced by Service Manager on page 681](#)
- [Specifying QoS Profiles in Service Definitions on page 674](#)
- [QoS for Service Manager Considerations on page 682](#)

Specifying QoS Profiles in Service Definitions

You can configure one QoS profile per subscriber interface. We recommend that you specify the QoS profile in the first set of services applied to the subscriber's interface.

You can modify the QoS profile by modifying configurations referenced by the QoS profile, including QoS parameter instances. You can also attach a new QoS profile when activating a service, but make sure that the QoS profile is attached to the subscriber's interface.

For more information about configuring QoS profiles, see the *Configuring and Attaching QoS Profiles to an Interface* chapter in the *JunosE Quality of Service Configuration Guide*.

To reference QoS profiles in a service definition, perform the following tasks:

- [Configuring a QoS Profile for Service Manager on page 674](#)
- [Specifying QoS Profiles in a Service Definition on page 675](#)

Configuring a QoS Profile for Service Manager

To configure a QoS profile for Service Manager:

1. Configure the profile.

```
host1(config)#profile videoService
host1(config-profile)#
```

When you specify the name of the profile to be configured, you enter Profile Configuration mode. You can specify a profile name with up to 80 alphanumeric characters.

Use the **no** version to remove a profile.

2. Configure the QoS profile.

```
host1(config-profile)#qos-profile videoBandwidth1
```

You can use this command to create a QoS profile on the router for use with Service Manager and enter QoS Profile Configuration mode. When the service is activated, the QoS profile is created and attached to the subscriber interface.

Use the **no** version to remove the QoS profile from the profile.

3. (Optional) Complete the QoS profile configuration described in the *Configuring and Attaching QoS Profiles to an Interface* chapter in the *JunosE Quality of Service Configuration Guide*.

Specifying QoS Profiles in a Service Definition

After you configure a QoS profile for Service Manager, you can reference it in a service definition. For example:

```
profile <# eastcoast ; '\n' #>
qos-profile <# video; '\n' #>
```

In this example, activating the service definition attaches the video QoS profile to the subscriber interface. Service Manager overwrites the existing QoS profile attachment at the subscriber interface.

Deactivating the service detaches the video QoS profile from the subscriber interface.

Related Documentation

- [Creating Service Definitions on page 668](#)
- [Managing Your Service Definitions on page 671](#)
- [Referencing QoS Configurations in Service Definitions on page 673](#)
- [Modifying QoS Configurations with Service Manager on page 677](#)
- [Removing QoS Configurations Referenced by Service Manager on page 681](#)
- *profile*
- *qos-profile*

Specifying QoS Parameter Instances in a Service Definition

You can specify that Service Manager create QoS parameter instances when the subscriber logs in (during service activation) or through RADIUS QoS parameter VSAs.

You can specify up to eight parameter instance commands within a profile. When you activate a service, Service Manager creates or modifies parameter instances for the subscriber interface that matches one of the subscriber-interface types configured in the QoS parameter definition.

Deactivating a service can modify or remove QoS parameter instances.

Using a service definition, you can also configure QoS parameters instances to add value to an existing parameter instance using the **add** keyword or dynamically create new parameter instances with an initial value using the **initial-value** keyword.

For more information about configuring QoS parameters, see *JunosE Quality of Service Configuration Guide*, *QoS Parameter Overview*.

To specify QoS parameter instances in a service definition, perform the following tasks:

- [Creating a Parameter Instance in a Profile on page 676](#)
- [Specifying QoS Parameter Instances in a Service Definition on page 676](#)

Creating a Parameter Instance in a Profile

To create a QoS parameter instance for Service Manager:

1. Configure the QoS parameter definition described in *JunosE Quality of Service Configuration Guide, QoS Parameter Overview*.

You must configure at least one controlled-interface type and one subscriber-interface type. The range specified in the parameter definition controls the available value of the parameter instance.

2. Configure the QoS profile.

```
host1(config)#profile video
```

3. Configure the QoS parameter instance command in the profile.

```
host1(config-profile)#qos-parameter videoBandwidth1 add 40000
```

When the service is activated, the parameter instances are created for the subscriber interface.

You can use the **add** keyword with this command in Profile Configuration mode to add a value to an existing parameter instance.

You can use the **initial-value** keyword with this command to create a new instance with the specified value.

```
host1(config-profile)#qos-parameter max-subscriber-bandwidth initial-value 15000
```

Use the **no** version to remove the QoS parameter instance command in the profile.

Specifying QoS Parameter Instances in a Service Definition

After you configure a QoS parameter instance for Service Manager, you can reference it in a service definition. For example:

```
<# qoserviceone(bandwidth1, bandwidth2) #>
profile <# profileName ; '\n' #>
qos-parameter <# qosParameterName1 ; ' ' ; bandwidth1 ; '\n' #>
qos-parameter <# qosParameterName2 ; ' ' ; bandwidth2 ; '\n' #>
<# endtmpl #>
```

When you activate a service, Service Manager creates the parameter instance and overwrites previous parameter instances. For example, activating the `qoserviceone` service definition creates a profile containing two QoS parameter instances. Service Manager creates the `qosParameterName1` parameter instance with the value of `bandwidth1`, and creates `qosParameterName2` with a value of `bandwidth2`.

If you activate the service definition using `qoserviceone(2000000,3000000)`, Service Manager creates `qosParameterName1` with a value of 2000000 and `qosParameterName2` instance with a value of 3000000.

You can use the **add** keyword to add value to an existing parameter instance. For example:

```
<# qoserviceone(bandwidth1, bandwidth2) #>
profile <# profileName ; '\n' #>
qos-parameter <# qosParameterName3 ; ' add ' ; bandwidth2 ; '\n' #>

<# endtmpl #>
```

When you specify parameter instances using the **add** keyword, you can also use the **initial-value** keyword to specify an initial value. For example:

```
<# qoserviceone(bandwidth1, bandwidth2) #>
profile <# profileName ; '\n' #>
qos-parameter <# qosParameterName4 ; ' add ' ; bandwidth2 ;
' initial-value 1000000' ; '\n' #>
<# endtmpl #>
```

When you activate the service, Service Manager locates the existing QoS parameter instance in the interface column. If Service Manager does not find a parameter instance, it creates one with a value specified in the **initial-value** keyword (in this case, 1000000). The value in the command is then added to the initial value. If an existing parameter instance is found, Service Manager adds the value to the existing interface.

For example, when you activate qosServiceOne as qosServiceOne(2000000,3000000), Service Manager attempts to locate the parameter instance qosParameterName4 for the subscriber's interface. If it finds a parameter instance, it adds bandwidth2 (3,000,000) to the current value. If Service Manager does not find a parameter instance, it creates one with an initial value of 1,000,000 and adds 3,000,000. The final parameter instance value is 4,000,000.

When deactivating the service, Service Manager locates the QoS parameter instance and subtracts the value in the command from the existing instance value. If the parameter is no longer referenced, the parameter instance is removed.

- Related Documentation**
- [Modifying QoS Configurations with Service Manager on page 677](#)
 - [Removing QoS Configurations Referenced by Service Manager on page 681](#)

Modifying QoS Configurations with Service Manager

This topic describes how to modify QoS configurations with Service Manager. To modify QoS configurations with Service Manager, perform the following steps:

- [Modifying Parameter Instances on page 677](#)
- [Modifying QoS Configurations in a Single Service Manager Event on page 679](#)
- [Modifying QoS Configurations by Using Other Sources on page 679](#)

Modifying Parameter Instances

Service Manager activates services without considering current parameter instance values. For example, when you deactivate a video service, Service Manager can add 5 Mbps to a parameter associated with the shaping rate of a video queue.

Similarly, Service Manager can deactivate services and restore parameter instances to their previous value. For example, when you deactivate a video service, Service Manager can subtract 5 Mbps from a parameter associated with the shaping rate of a video queue.

[Table 151 on page 678](#) lists the results of a series of activations and deactivations of parameters using the **add** and **initial-value** keywords.

Table 151: Sample Modifications Using the Add and Initial-Value Keywords

Action	QoS Parameter Instance	Result
Activate	qos-parameter video-bw add 5000000 initial-value 0	Parameter instance video-bw is created with a value of 5000000
Activate	qos-parameter video-bw add 1000000 initial-value 0	Parameter instance video-bw is increased by 1000000, for a total of 6000000
Deactivate	qos-parameter video-bw add 1000000 initial-value 0	Parameter instance video-bw is decreased by 1000000, for a total of 500000
Deactivate	qos-parameter video-bw add 5000000 initial-value 0	Parameter instance video-bw is removed

Removing a parameter instance using profiles is based on the number of times a parameter instance is modified, not the value added.

Modifying parameter instances in profiles and modifying explicit parameter instances can cause invalid parameter instance values. [Table 152 on page 678](#) lists a series of activations and deactivations using parameter instances in profiles and explicit parameter instances. By the second deactivation, the parameter has a negative value (-4000000).



NOTE: We recommend that you do not configure negative values for Service Manager.

Table 152: Sample Modifications Using Parameter Instances

Action	QoS Parameter Instance	Result
Activate	qos-parameter video-bw add 5000000 initial-value 0	Parameter instance video-bw is created with a value of 5000000
Activate	qos-parameter video-bw add 1000000 initial-value 0	1000000 is added to parameter instance video-bw, for a total of 6000000
Activate	qos-parameter video-bw 2000000	Parameter instance video-bw is set to 2000000

Table 152: Sample Modifications Using Parameter Instances (*continued*)

Action	QoS Parameter Instance	Result
Deactivate	qos-parameter video-bw add 1000000 initial-value 0	1000000 is subtracted from parameter instance video-bw for a total of 1000000
Deactivate	qos-parameter video-bw add 5000000 initial-value 0	5000000 is subtracted from parameter instance video-bw for a total of -4000000
Deactivate	qos-parameter video-bw 2000000	Parameter instance video-bw is removed

Modifying QoS Configurations in a Single Service Manager Event

QoS accepts QoS profile attachments and parameter instances created using multiple sources (profiles, RADIUS, or Service Manager) within a single Service Manager event. Events include:

- Subscriber login
- Subscriber logout
- RADIUS Change of Authority

QoS prioritizes the creation of QoS profiles and parameter instances within events.

[Table 153 on page 679](#) lists the sources that overwrite QoS profiles and parameter instances created by other sources. Each row represents new QoS profiles and parameter instances; columns represent existing QoS profiles and parameter instances.

Table 153: Configuration Within a Single Service Manager Event

	Profile	RADIUS	Service Manager
Profile	✓	–	–
RADIUS	✓	✓	–
Service Manager	✓	✓	✓

Modifying QoS Configurations by Using Other Sources

You can modify QoS configurations with Service Manager by using other QoS sources. For example, you can modify a parameter instance that was created with Service Manager by using the CLI. Similarly, you can use SNMP to detach a QoS profile attached by Service Manager.

[Table 154 on page 680](#) lists the sources that you can use to modify QoS profile attachments and parameter instances.

Table 154: Modifying QoS Configurations with Other Sources

	QoS Profile Attachment	QoS Parameter Instances
Service Manager	✓	✓
RADIUS	✓	✓
SNMP	✓	–
SRC software	✓	–
CLI	✓	✓

The following sections describe the precedence of each source when modifying configurations.

Service Manager

QoS profile attachments and parameter instances created through Service Manager have precedence over all other sources. For example, Service Manager can overwrite a QoS profile attachment modified through RADIUS, SNMP, the SRC software, or the CLI.

Conversely, Service Manager configurations can be overwritten through SNMP, the SRC software, and the CLI, but not by RADIUS.

Service Manager counts references of parameter instances. You can modify parameter instances created by Service Manager using other sources without affecting the reference counts. For more information, see [“QoS for Service Manager Considerations” on page 682](#).

RADIUS

QoS profile attachments and parameter instances configured through RADIUS can overwrite QoS profile attachments and parameter instances configured through the SNMP, the SRC software, and the CLI, but not those created by Service Manager.

Conversely, QoS profiles and parameter instances configured through RADIUS can be overwritten by any source (SNMP, the SRC software, CLI, and Service Manager).

SNMP, the SRC Software, and the CLI

QoS profile attachments and parameter instances configured through the CLI can overwrite QoS profile attachments and parameter instances configured through any source.

QoS profiles attached through SNMP and the SRC software can also overwrite QoS profile attachments configured through any source.

Conversely, QoS profiles and parameter instances configured through the CLI, SNMP, or the SRC software can be overwritten by any source.

Related Documentation


- [Creating Service Definitions on page 668](#)
- [Managing Your Service Definitions on page 671](#)
- [Specifying QoS Profiles in Service Definitions on page 674](#)
- [Referencing QoS Configurations in Service Definitions on page 673](#)
- [Removing QoS Configurations Referenced by Service Manager on page 681](#)
- [QoS for Service Manager Considerations on page 682](#)

Removing QoS Configurations Referenced by Service Manager

When Service Manager no longer references a QoS configuration, it must be removed from the service definition.

Figure 30 on page 681 shows the references for QoS configurations.

Figure 30: QoS Configuration Dependency Chain



Service Manager automatically removes QoS profiles and parameter instances. After removing the QoS profile and parameter instances, Service Manager automatically removes the following QoS configurations in the following order:

1. QoS profiles
2. Scheduler profiles
3. Queue profiles
4. Drop profiles
5. Statistics profiles

Service Manager does not automatically remove the following QoS configurations:

- Parameter definitions
- Traffic classes

- Traffic-class groups
- QoS-port-type profiles

Related Documentation

- [Creating Service Definitions on page 668](#)
- [Managing Your Service Definitions on page 671](#)
- [Specifying QoS Profiles in Service Definitions on page 674](#)
- [Referencing QoS Configurations in Service Definitions on page 673](#)
- [Modifying QoS Configurations with Service Manager on page 677](#)
- [QoS for Service Manager Considerations on page 682](#)

QoS for Service Manager Considerations

When you specify QoS configurations in Service Manager, the following considerations apply.

RADIUS or Service Manager

We recommend that you choose either RADIUS or Service Manager to create a single parameter instance. If you use both RADIUS and Service Manager, parameter instances activated using Service Manager take precedence.

Interoperability with Other Service Components

Service Manager removes QoS profiles and parameter instances if other components in the service definition (for example, policies) cause an error.

QoS Statistics

Service Manager counts references of parameter instances in profiles. The reference count is incremented each time the parameter is configured through the CLI, RADIUS, or Service Manager. The reference count is decremented each time the parameter is unconfigured, such as through service deactivation. Modifications to parameter instances are also reference counted, using a separate reference count. Parameter instances are removed when both reference counts reach zero.

Service Manager also counts references of modified parameters in profiles using the **add** keyword. The reference count is incremented each time the parameter is modified through service activation with the **add** keyword. The reference count is decremented each time the parameter is modified through service deactivation. References of regular parameter instances are also counted, using a separate reference count. Parameter instances are removed when both reference counts reach zero.

Ranges

You can verify ranges for parameter instances by specifying a range in the parameter definition using the **range** command.

When activating the service or modifying parameters, Service Manager verifies the value of the parameter instance to be within the specified range and generates an informational log message indicating the value is outside the range. Service Manager does not verify ranges when you specify the parameter instances within profiles at the time of configuration.

Related Documentation

- [Creating Service Definitions on page 668](#)
- [Managing Your Service Definitions on page 671](#)
- [Specifying QoS Profiles in Service Definitions on page 674](#)
- [Referencing QoS Configurations in Service Definitions on page 673](#)
- [Modifying QoS Configurations with Service Manager on page 677](#)
- [Removing QoS Configurations Referenced by Service Manager on page 681](#)

Configuring the Service Manager License

Use the Service Manager license to enable full Service Manager application support. You can create a maximum of 10 subscriber sessions when the Service Manager license is not enabled. If you disable the Service Manager license and more than 10 subscriber sessions exist, you cannot enable any new sessions—however, all existing active subscriber sessions continue to function.

For information about the maximum number of subscriber sessions supported, see *JunosE Release Notes, Appendix A, System Maximums*.

To specify the Service Manager license and enable full Service Manager application support—if the license is not enabled, you are limited to 10 subscriber sessions:



NOTE: Obtain the license from Juniper Networks Customer Service or your Juniper Networks sales representative.

- Issue the **license service-management** command in Global Configuration mode.

```
host1(config)#license service-management 123456789
```

The license is a unique string of up to 15 alphanumeric characters.

Use the **no** version to disable the license.

Related Documentation

- [Configuring Service Manager on page 665](#)
- [Managing Your Service Definitions on page 671](#)
- [license service-management](#)

Managing and Activating Service Sessions Overview

You can use either RADIUS or the CLI to manage, activate, and deactivate service sessions. The following list describes some of the differences between using RADIUS and the CLI to manage the Service Manager application.

- RADIUS-based login and RADIUS COA support:
 - Provides dynamic activation and deactivation based on subscriber service selection
 - Provides greater flexibility and efficient management for a large number of subscribers and services
 - Enables you to use mutual exclusion (mutex) groups to create mutex services (RADIUS COA only)
- CLI-based support:
 - Provides static activation and deactivation for subscribers who are always logged in
 - Is useful for testing new service definitions
 - Enables you to preprovision services that you can activate later

Related Documentation

- [Service Session Profiles Overview on page 711](#)
- [Understanding Service Manager RADIUS Attributes on page 686](#)
- [Managing Subscriber Service Sessions by Using the CLI Overview on page 707](#)
- [Activating Subscriber Service Sessions by Using the CLI Overview on page 708](#)
- [Deactivating Subscriber Service Sessions by Using the CLI Overview on page 715](#)

Managing Subscriber Service Sessions by Using RADIUS Overview

Service Manager supports two RADIUS-based methods for dynamically activating subscriber service sessions. Dynamic service sessions that RADIUS activates are not stored in NVS. Both methods can also apply optional statistics and session threshold (volume and time) configurations. The two methods differ in how Service Manager activates a subscriber service session:


- RADIUS login method—The service session is activated when the subscriber logs in. At login, RADIUS verifies that the Activate-Service attribute is configured in the subscriber's RADIUS record. RADIUS then uses vendor-specific attributes (VSAs) in the Access-Accept packet to activate the service session for the subscriber. This method is useful when your subscribers are not currently logged in.
- RADIUS COA method—Supports dynamic service selection for subscribers. For example, the subscriber might have logged in without a service, or might have used the RADIUS login method to activate a service at login. If no service was activated at login (because of no Activate-Service attribute in the user's RADIUS record), you can later use the

COA method and a separate RADIUS record to create a subscriber session and activate a service session for the subscriber. Or, if the RADIUS login method was used and the subscriber already has an active service session, you can use the COA method and a new RADIUS record to activate a new service session for the subscriber (and optionally deactivate the existing service session). The RADIUS COA method is useful when you have a large number of users already logged in through RADIUS and you want to activate new services for them. This method is also used for the guided entrance service described in “[Example: Guided Entrance Service Definition](#)” on page 731.

The RADIUS COA method also supports the use of mutex groups to create mutex services. See “[Activating and Deactivating Subscriber Services by Using Mutex Groups Overview](#)” on page 691.

Figure 31 on page 685 compares the two RADIUS-based methods.

Figure 31: Comparing RADIUS Login and RADIUS COA Methods



Related Documentation

- [Service Session Profiles Overview](#) on page 711
- [Managing and Activating Service Sessions Overview](#) on page 684
- [Activating Subscriber Service Sessions Using RADIUS](#) on page 685
- [Deactivating Service Sessions Using RADIUS](#) on page 690
- [Activating and Deactivating Subscriber Services by Using Mutex Groups Overview](#) on page 691

Activating Subscriber Service Sessions Using RADIUS

To use RADIUS to activate subscriber service sessions, you create a RADIUS record that includes the Activate-Service VSA. For the RADIUS login method, this RADIUS record is used by the Access-Accept message to start Service Manager and activate the service when the subscriber logs in.

For the RADIUS COA method, the service provider uses a CoA-Request message to activate and deactivate the service for the subscriber who is already logged in.

To configure a service session that will be activated by RADIUS:

1. Create the RADIUS record for the subscriber and service:

- For RADIUS login—Create the RADIUS record for the subscriber and include the Activate-Service VSA in the record. Specify values for the parameters defined in the service template name of the definition macro file.
- For RADIUS COA—Format the COA message to create the RADIUS record for the subscriber. Include the Activate-Service VSA in the record. Optionally, include the Deactivate-Service VSA if the subscriber has an active service session that you want to deactivate. Specify values for the parameters defined in the service template name of the definition macro file.



NOTE: You specify the parameter values in the order in which the parameters appear in the template name of the service definition file. For example, in the tiered service that is defined in the sample service definition macro file in [“Creating Service Definitions” on page 668](#), the template name is:

```
<# tiered(inputBW, outputBW) #>
```

For the RADIUS Activate-Service VSA, you specify values for the input and output bandwidth:

```
tiered(1280000, 5120000)
```

2. Specify optional VSAs for the service session as needed:

- Service-Volume
- Service-Timeout
- Service-Statistics

**Related
Documentation**

- [Service Session Profiles Overview on page 711](#)
- [Understanding Service Manager RADIUS Attributes on page 686](#)
- [Deactivating Service Sessions Using RADIUS on page 690](#)
- [Activating and Deactivating Subscriber Services by Using Mutex Groups Overview on page 691](#)
- [Activating and Deactivating Multiple Services on page 692](#)

Understanding Service Manager RADIUS Attributes

For the RADIUS login method, the RADIUS VSAs for service activation, threshold configuration, statistics configuration, and interim accounting in Access-Accept messages at subscriber login are used by Service Manager to activate the appropriate service session. For the RADIUS COA method, Service Manager uses the VSAs for service activation and deactivation, threshold configuration, statistics configuration, and interim accounting in COA-Request messages to activate the service session. The accounting-related VSAs are included in RADIUS accounting messages.

Table 155 on page 687 lists the Service Manager-related attributes and indicates which are tagged VSAs. See “Using Tags with RADIUS Attributes” on page 689 for a discussion about using tagged VSAs to group attributes for a service.

Table 155: Service Manager RADIUS Attributes

Attribute Number	Attribute Name	RADIUS Message Type	VSA Description
[1]	User-Name (used with Virtual-Router, Juniper Networks VSA 26-1)	Access-Accept	Uniquely identifies the subscriber session
[8]	Framed-IP-Address (used with Virtual-Router, Juniper Networks VSA 26-1)	Access-Accept	Uniquely identifies the subscriber session
[26-65]	Activate-Service	Access-Accept and COA-Request	Name of the service to be activated; includes parameter values; a tagged VSA
[26-66]	Deactivate-Service	Access-Accept and COA-Request	Name of the service to be deactivated NOTE: This VSA is only used by COA.
[26-67]	Service-Volume	Access-Accept and COA-Request	Number of MB of traffic that the service can consume; the service is terminated when output byte count exceeds this value; a tagged VSA
[26-68]	Service-Timeout	Access-Accept and COA-Request	Number of seconds that the service is to remain active; the service is terminated when the time expires; a tagged VSA
[26-69]	Service-Statistics	Access-Accept and COA-Request	Statistics configuration; a tagged VSA: 0 = disable 1 = timestamp only 2 = timestamp and volume
[26-83]	Service-Session	For service sessions only: Acct-Start Acct-Stop Interim-Acct	Name of the service (including parameter values) with which the statistics are associated
[26-140]	Service-Interim-Acct-Interval	Access-Accept and COA-Request	Number of seconds between accounting updates for a service; a tagged VSA
[31]	Calling-Station-ID	Access-Accept	Uniquely identifies the subscriber session

Table 155: Service Manager RADIUS Attributes (*continued*)

Attribute Number	Attribute Name	RADIUS Message Type	VSA Description
[44]	Acct-Session-ID	Acct-Start Acct-Stop Interim-Acct	Accounting identifier that makes it easy to match start and stop records in a log file; the format is extended to include a colon-separated value that uniquely identifies the subscriber session



NOTE: Service Manager statistics collection is a three-part procedure. You must configure statistics information in the service definition macro file, enable statistics collection in the RADIUS record, and also enable statistics collection for the policy referenced in the service macro using the **statistics enabled** keyword in the command used for policy attachment in the profile.

The Service-Volume and Service-Timeout VSAs rely on the values captured by the Service Manager statistics feature to determine when a threshold is exceeded. Therefore, you must configure and enable statistics collection to use these attributes. Service-Volume For detailed information about Service Manager statistics see [“Configuring Service Manager Statistics” on page 718](#).

Table 156 on page 688 describes a partial RADIUS Access-Accept packet that activates a service session for subscriber client1@isp1.com. (The figure in [“Creating Service Definitions” on page 668](#) shows the service definition macro file that creates the tiered service.) The session enables the subscriber to use the tiered service with an input bandwidth of 1280000 and output bandwidth of 5120000. The subscriber can use the service for 5 hours (18000 seconds), and Service Manager captures both timestamp and volume statistics during the session (service-statistics value of 2). Also, accounting for the service is updated every 600 seconds (10 minutes).

Table 156: Sample RADIUS Access-Accept Packet

RADIUS Attribute	Tag	Value
username	none	client1@isp1.com
class	none	(binary data)
service-activation	6	tiered(1280000, 5120000)
service-timeout	6	18000
service-statistics	6	2
service-interim-acct-interval	6	600

Using Tags with RADIUS Attributes

Service Manager uses tagged RADIUS VSAs to enable a single RADIUS record to activate multiple service sessions for a subscriber, with each session having unique attributes. A particular tag identifies a specific Activate-Service attribute and all other RADIUS attributes that are associated with that Activate-Service attribute.

You can specify a maximum of 8 tags (1–8), which enables you to activate up to eight unique service sessions for a subscriber in a single RADIUS record. The following are tagged VSAs—they must always have a tag in their RADIUS entry:

- Activate-Service
- Service-Statistics
- Service-Timeout
- Service-Volume
- Service-Interim-Acct-Interval

Table 157 on page 689 describes an Access-Accept packet that activates the two services, tiered and voice, for subscriber client1@isp1.com. Each service has its own unique tag, enabling you to assign attributes for one service, but not the other. For example, the two services have different timeout settings and different interim accounting intervals, and statistics are enabled only for the tiered service.

Table 157: Using Tags

RADIUS Attribute	Tag	Value
username	none	client1@isp1.com
class	none	(binary data)
service-activation	2	tiered(1280000, 5120000)
service-timeout	2	18000
service-statistics	2	1
service-interim-acct-interval	2	600
service-activation	6	voice(100000)
service-timeout	6	1440
service-interim-acct-interval	6	1200

Related Documentation

- [Service Session Profiles Overview on page 711](#)
- [Working with Service Session Profiles on page 712](#)

- [Managing and Activating Service Sessions Overview on page 684](#)
- [Managing Subscriber Service Sessions by Using RADIUS Overview on page 684](#)
- [Understanding RADIUS Accounting for Service Manager on page 697](#)
- [Activating Subscriber Service Sessions Using RADIUS on page 685](#)
- [Deactivating Service Sessions Using RADIUS on page 690](#)

Deactivating Service Sessions Using RADIUS

You can deactivate a service session by using a COA-Request message or when a subscriber logs out of a RADIUS-activated service session. If the subscriber logs off the router, Service Manager deactivates that subscriber session and all associated service sessions.

RADIUS also supports attributes that you can use to manage deactivation of service sessions. You can deactivate service sessions with the following set of tasks:

- [Setting Time or Volume Thresholds for a Service on page 690](#)
- [Using the Deactivate-Service Attribute on page 691](#)

Setting Time or Volume Thresholds for a Service

You can set a threshold for the session by including one or both of the following attributes in the RADIUS record:



NOTE: The Service-Timeout and Service-Volume attributes use values captured by the Service Manager statistics feature to determine when a threshold is exceeded. Therefore, you must configure and enable statistics collection to use these attributes. See [“Configuring Service Manager Statistics” on page 718](#).

- **Service-Timeout**—The number of seconds that the service session is active. You can specify a number in the range 0–16777215. Values greater than 16777215 are recycled, starting from the initial value of 0. For example, if you specify the value for Service-Timeout VSA as 16777218, this value is equivalent to 2 for this VSA. A value of 0 indicates that the session never times out. A particular Service-Timeout VSA can be used by a maximum of 2000 services.

The service-timeout threshold accuracy is within 30 seconds of the specified value.

- **Service-Volume**—The total number of MB of traffic that can use the service session. You can specify a number in the range 0–16777215 MB. Values greater than 16777215 are recycled, starting from the initial value of 0. A value of 0 indicates that there is no limit to the amount of traffic for the session. For example, if you specify the value for Service-Timeout and Service-Volume VSAs as 16777216 and 16777217, these values are equivalent to 0 and 1 respectively for these VSAs. A particular Service-Volume VSA can be used by a maximum of 1000 services.



NOTE: Service Manager terminates a session when the *output* byte count exceeds the configured service-volume threshold. The output byte count is captured by the `output-stat-clacl` string in the classifier list variable that you configure to collect statistics. See [“Configuring Service Manager Statistics” on page 718](#).

The service-volume threshold accuracy is based on a 10-second period. Service Manager does not immediately deactivate a service session when the output byte count reaches the service-volume threshold. Instead, Service Manager checks the volume in 10-second intervals and deactivates a service session at the end of the 10-second period in which the output byte count reaches the volume threshold. For example, if a threshold is reached 4 seconds into the 10-second interval, the session continues for the remaining 6 seconds in the measuring period and is then terminated. Therefore, the total volume equals the threshold plus the volume during the additional 6 seconds.

When the output byte count reaches the threshold, RADIUS deactivates the service session. You must use tags to associate threshold attributes with the Activate-Service attribute for the service session.

Using the Deactivate-Service Attribute

You can also include the Deactivate-Service attribute in the subscriber's RADIUS record. The format for this attribute is the same as the format of the Activate-Service attribute—the name of the service, including parameters. The Deactivate-Service attribute is used by RADIUS COA messages, such as in a guided entrance service. See [“Example: Guided Entrance Service Definition” on page 731](#) for more information.

Related Documentation

- [Service Session Profiles Overview on page 711](#)
- [Understanding Service Manager RADIUS Attributes on page 686](#)
- [Understanding RADIUS Accounting for Service Manager on page 697](#)
- [Activating Subscriber Service Sessions Using RADIUS on page 685](#)
- [Activating and Deactivating Subscriber Services by Using Mutex Groups Overview on page 691](#)
- [Activating and Deactivating Multiple Services on page 692](#)

Activating and Deactivating Subscriber Services by Using Mutex Groups Overview

Service Manager supports two methods that use RADIUS COA-Request messages to activate and deactivate subscriber services and that can also dynamically change a service that is currently provided to a subscriber.

In the first method, you use a COA message with the Activate-Service VSA to activate the new service; you can optionally include the Deactivate-Service VSA to deactivate the subscriber's existing service. This method is described in [“Activating Subscriber Service Sessions Using RADIUS” on page 685](#).

The second method uses mutual exclusion (mutex) groups to create mutex services. With this method, you group services together in a mutex group. When you use a COA message to activate a service that is in a mutex group, Service Manager activates the new service and implicitly deactivates any existing service that it is a member of the same mutex group as the newly activated service. Service Manager does not deactivate an existing service that is a member of a different mutex group or is not a member of a mutex group.

Using mutex services results in a more reliable activation and deactivation process than the original COA-Request method. With mutex services, Service Manager always activates the new service before deactivating the existing service. This ensures that the subscriber is never without an active service. In the original COA-Request method, the order of activation and deactivation is random—in some cases the existing service might be deactivated before the new service is activated, or the new activation might fail. In these cases, the subscriber might be without an active service.

If statistics are enabled when you activate a mutex service, Service Manager sends a RADIUS Acct-Stop message for the deactivated service.

**Related
Documentation**

- [Service Session Profiles Overview on page 711](#)
- [Understanding Service Manager RADIUS Attributes on page 686](#)
- [Activating Subscriber Service Sessions Using RADIUS on page 685](#)
- [Deactivating Service Sessions Using RADIUS on page 690](#)
- [Configuring a Mutex Service on page 693](#)
- [Activating and Deactivating Multiple Services on page 692](#)

Activating and Deactivating Multiple Services

The Service Manager mutex service feature enables you to activate and deactivate multiple services with a single COA-Request message. A COA-Request message can have more than one service activation request—the multiple service requests might be from the same mutex group or from different groups. The following examples describe how you might use mutex groups to activate and deactivate multiple services.

- Case 1—Multiple mutex services of the same mutex group

Service Manager activates the multiple mutex services, which are in the same group, then deactivates all previously existing services that are also members of that mutex group. Active services that are members of different mutex groups are unaffected.

- Case 2—Multiple mutex services of different mutex groups

Service Manager activates the mutex services, which are members of different mutex groups. Service Manager then deactivates all previously existing services that are members of the same mutex groups as any of the newly activated services. Active services that are members of different mutex groups are unaffected.

Related Documentation

- [Service Session Profiles Overview on page 711](#)
- [Understanding Service Manager RADIUS Attributes on page 686](#)
- [Activating Subscriber Service Sessions Using RADIUS on page 685](#)
- [Deactivating Service Sessions Using RADIUS on page 690](#)
- [Activating and Deactivating Subscriber Services by Using Mutex Groups Overview on page 691](#)
- [Configuring a Mutex Service on page 693](#)

Configuring a Mutex Service

To configure and enable a mutex service, you complete the following steps:

1. Create the new service definition and configure the service as a member of a mutex group.

When you create the service definition, include the following service attribute in the service definition, where `groupIndex` identifies the mutex group for this service. The `groupIndex` can be a number in the range -1 to -2147483647 or 1 to 2147483646. If the `groupIndex` is outside of the acceptable ranges, or if you do not include the `mutex-group` statement, the service is not included in a mutex group.

```
<# env.setResult("mutex-group", "groupIndex" ) #>
```

For example (the mutex group attribute is highlighted in bold text):

```
!parameterizes input and output bandwidth
<# tiered(inputBW, outputBW) #>

<# uid := app.servicemanager.getUniqueId #>
<# name := "SM-tiered-" $ uid #>
<# oname := "SM-O-tiered-" $ uid #>

classifier-list matchAll ip any any
rate-limit-profile <# name #> one-rate
committed-rate <# inputBW; '\n' #>

policy-list <# name; '\n' #>
classifier-group matchAll precedence 10000
rate-limit-profile <# name; '\n' #>
traffic-class best-effort

policy-list <# oname; '\n' #>
classifier-group matchAll precedence 10000
traffic-class best-effort

profile <# name; '\n' #>
ip policy secondary-input <# name #> statistics enabled merge
ip policy output <# oname #> statistics enabled merge
qos-profile triplePlayIP
qos-parameter maxSubscBW <# outputBW; '\n' #>

<# env.setResult("mutex-group", "12" ) #>
<# env.setResult("activate-profile", name) #>
<# env.setResult("secondary-input-stat-clac1", "matchAll") #>
<# env.setResult("output-stat-clac1", "matchAll") #>
```

```
<# endtmp1 #>
```

2. Activate the mutex service

Use a RADIUS COA-Request message and the new service definition to create the mutex service. The new service is considered a mutex service because it belongs to a mutex group.

Service Manager activates the new service and deactivates any existing active service that is a member of the same mutex group as the new service.

3. (Optional) Verify the status of the new service.

```
host1# show service-management subscriber-session client1@isp.com interface ip 192.168.0.1
User Name: CLIENT1@ISP.COM, Interface: ip 192.168.0.1
Id: 1
Owner: AAA 4194326
Non-volatile: False
State: Active
ServiceSessions:
  Name                mutex  Owner/Id      State                Operation
  -----
  tiered(2000000,3000000) 12    AAA 4194326  ConfigApplySuccess  Activate
  Name                Non-volatile
  -----
  tiered(2000000,3000000) False
```

Related Documentation

- [Service Session Profiles Overview on page 711](#)
- [Understanding Service Manager RADIUS Attributes on page 686](#)
- [Activating Subscriber Service Sessions Using RADIUS on page 685](#)
- [Deactivating Service Sessions Using RADIUS on page 690](#)
- [Activating and Deactivating Subscriber Services by Using Mutex Groups Overview on page 691](#)
- [Activating and Deactivating Multiple Services on page 692](#)

Combined and Independent IPv4 and IPv6 Services in a Dual Stack Overview

Internet Protocol version 6 (IPv6) is designed to enhance IP addressing and maintain other IPv4 functions that work well. Organizations worldwide are developing new applications to take advantage of the many feature enhancements within IPv6 to help bring back end-to-end controlled communications across a transparent network infrastructure. To ensure optimum performance, such applications implement a dual-stack architecture in which IPv4 and IPv6 protocols share a common transport and framing layer.

A dual-stack implementation supports both IPv4 and IPv6 hosts to help provide a smooth transition to all parts of an enterprise network. With this flexible method of implementation, providers can carry IPv6 traffic over their existing core networks and customers can roll out IPv6 to more sites.

The PPP link between the customer premises equipment (CPE) and the provider edge (PE) device or E Series router equipment might require both IPv4 and IPv6 protocols for transmission of data. Such networks require that PE devices run a dual stack of IPv4 and IPv6 services. In this release, the Service Manager application on the E Series router supports authentication, service selection, and service activation and deactivation to subscribers for both IPv4 and IPv6 protocols in a dual stack configuration.

You can configure services in a dual stack for IPv4 and IPv6 either independently or as a single entity. Service Manager only tracks JunosE objects that are passed back in the `env.setResult` method when a service definition is executed. In an IPv6 environment, you must modify the service definition macro file to include the objects that the Service Manager requires to categorize a service as IPv4, IPv6, or a combination of both IPv4 and IPv6.

You can use the **service-interface-type** object in the service definition macro file to specify whether a service must be defined for IPv4 or IPv6. Configuring the `service-interface-type` object is not mandatory if a service is required only for IPv4 or L2TP subscribers. However, you must specify the `service-interface-type` object when a service is required for IPv6 subscribers or IPv4 and IPv6 subscribers in a dual stack. After you create a new service definition file with the `service-interface-type` object and install it on the router, the Service Manager determines whether a service must be tagged as IPv4, IPv6, or a combination of the two by parsing the objects passed using the macro environment command, **env.setResult**. The `service-interface-type` object can be configured with one of the following values:

- `ipv4`—Specifies that the service session profile must be applied to IPv4 interfaces only. This object is optional only when IPv4 or L2TP subscribers are in a network.
- `l2tp`—Specifies that the service session profile must be applied to L2TP interfaces. This object is optional only when IPv4 or L2TP subscribers are in a network.
- `ipv4-ipv6`—Specifies that the service session profile must be applied to both IPv4 and IPv6 interfaces in a dual stack. You must configure this object when IPv6 subscribers or IPv4/IPv6 subscribers in a dual stack are in a network.
- `ipv6`—Specifies that the service session profile must be applied to IPv6 interfaces only. You must configure this object when IPv6 subscribers or IPv4 and IPv6 subscribers in a dual stack are in a network.

When you create the service definition, include the following service attribute in the service definition if you want the service to be defined for IPv4 interfaces only. The profile identifier returned from the `activate-profile` object is applied to IPv4 interfaces.

```
<# env.setResult("service-interface-type", "ipv4" ) #>
```

To configure a service macro to be used for IPv6 interfaces only, specify the following object in the macro definition file. The profile identifier returned from the `activate-profile` object is applied to IPv6 interfaces.

```
<# env.setResult("service-interface-type", "ipv6" ) #>
```

To configure a service macro to be used for IPv4 and IPv6 interfaces in a dual stack, specify the following object in the macro definition file. The profile identifier returned from the activate-profile object is applied to both IPv4 and IPv6 interfaces.

```
<# env.setResult("service-interface-type", "ipv4-ipv6" ) #>
```

**Related
Documentation**

- [Service Definitions Overview on page 667](#)
- [Activation and Deactivation of IPv4 and IPv6 Services in a Dual Stack on page 696](#)
- [Service Interim Accounting for IPv4 and IPv6 Services in a Dual Stack Overview on page 705](#)
- [Example: Combined IPv4 and IPv6 Service in a Dual Stack Service Definition on page 740](#)

Activation and Deactivation of IPv4 and IPv6 Services in a Dual Stack

You can configure IPv4 and IPv6 services in a dual stack either as independent services or as a combined service. The following sections describe the two types of configurations and their behavior when they are activated and deactivated.

Independent IPv4 and IPv6 Services in a Dual Stack

To configure separate services for IPv4 and IPv6 interfaces you must create and install separate service definitions on the router. For example, you can create a service definition called `iponeV4` to be used for IPv4 traffic and service definition called `iponeV6` to be used for IPv6 traffic. Both the services defined for IPv4 and IPv6 must be configured for the subscriber on the RADIUS server. When the subscriber is authenticated using RADIUS authentication, two services, one each for IPv4 and IPv6, are created. The subscriber service sessions are created and activated when the subscriber logs in using the RADIUS Access-Accept messages, Change-of-Authorization-Request (COA-Request) messages, or CLI commands. After the subscriber service session is activated, the policies defined in the interface profile specified by the activate-profile object in the service macro file are applied to the IPv4 and IPv6 interfaces. Service session profiles provide additional flexibility to the Service Manager application by enabling you to assign one or more supported attributes to a particular activation of a service.

Deactivation of service sessions is also performed for each individual service. If an interface is deleted, all the services associated with that interface are also deleted. For example, if you delete an IPv6 interface, all the services associated with IPv6 are deleted. However, IPv4 subscriber service sessions are not disrupted. When a user logs out of a session, all services associated with the subscriber session are also removed along with the subscriber session. If the subscriber has two services and one of them was not successfully applied to an interface, then that service is removed. For example, if a subscriber has two services, `iponev4` and `iponev6`, configured on the RADIUS server and only one service was successfully configured on an interface, then the failed service is deleted when the service is deactivated.

Combined IPv4 and IPv6 Service in a Dual Stack

To configure a single service for IPv4 and IPv6 interfaces, you can create and install one service definition on the router that handles the traffic for both these protocols. For

example, you can create a service definition called `iponeV4V6` to be used for both IPv4 and IPv6 traffic. This service must be configured for the subscriber on the RADIUS server. When the subscriber is authenticated using RADIUS authentication, a single service is created and activated using the RADIUS or CLI client type that Service Manager supports. After the subscriber service session is activated, the policies defined in the interface profile specified by the `activate-profile` object in the service macro file are applied to both IPv4 and IPv6 interfaces. The elements in the profile to be attached to the interfaces are determined by the type of the interface. The combined service session is active if either of the two conditions is satisfied:

- Both the IPv4 and IPv6 interfaces are up
- Either the IPv4 or IPv6 interface is up

Deactivation of service sessions results in disconnection of services for both IPv4 and IPv6 subscribers.

Performance Impact on the Router and Compatibility with Previous Releases for an IPv4 and IPv6 Dual Stack

In an environment where only IPv4 subscribers exist, the memory usage on the E Series router is the same as the usage in previous releases. If IPv4 and IPv6 are configured as independent services, memory usage increases because each IPv6 service is counted as a separate service and uses all the system resources than an IPv4 service requires. Memory impact in such a case is proportional to the total number of services configured. You can view the number of service sessions currently active for a subscriber by viewing the Service Sessions field from the output of the **show service-management** command.

If you configured a combined IP4 and IPv6 service, the memory usage is the same as that required for one subscriber service session. The number of subscribers that are supported by the line modules depends on the number of available resources, such as external parent groups. If you configure services that are to be used in an IPv4 or L2TP network, you need not change the previously defined service macros. However, if a subscriber requires the service macro applied to IPv6 interfaces or wants to apply a combined policy for both IPv4 and IPv6 interfaces, you must modify the macro file for the appropriate service interface type.

Related Documentation

- [Service Definitions Overview on page 667](#)
- [Creating Service Definitions on page 668](#)
- [Combined and Independent IPv4 and IPv6 Services in a Dual Stack Overview on page 694](#)
- [Example: Combined IPv4 and IPv6 Service in a Dual Stack Service Definition on page 740](#)

Understanding RADIUS Accounting for Service Manager

The Service Manager application supports RADIUS accounting and interim accounting for subscriber service sessions that are activated by the RADIUS login and RADIUS COA methods. When RADIUS accounting is enabled, RADIUS generates:

- An Acct-Start message when a service session is activated
- An Acct-Stop message when a service session is deactivated
- Interim-Acct messages

RADIUS accounting messages always include Service Manager time statistics. You must enable Service Manager volume statistics for a service session.

When you terminate a subscriber session, Service Manager first sends RADIUS Acct-Stop messages for any active services associated with the subscriber session, and then sends the Acct-Stop message for the subscriber session.



NOTE: Service Manager statistics collection is a three-part procedure. You must configure statistics information in the service definition macro file, enable statistics collection by either RADIUS or the CLI, and also enable statistics collection for the policy referenced in the service macro using the **statistics enabled** keyword in the command for policy assignment to a profile at the time of attachment of the policy to an interface. For detailed information about Service Manager statistics, see [“Configuring Service Manager Statistics” on page 718](#).

To support RADIUS accounting for Service Manager, the RADIUS Acct-Session-ID attribute [44] has been extended to include a colon-separated identifier, which uniquely identifies a service for a subscriber. For example:

```
erx FastEthernet 12/0:0001048580:002478
```

The Service-Session attribute (VSA 26-83) identifies the name of the service. This attribute is the value of the Activate-Service or Deactivate-Service attribute (including parameter values) that was used in the RADIUS Access-Accept message to activate or deactivate the service session. For example:

```
tiered(1280000, 5120000)
```

[Table 158 on page 698](#) lists the RADIUS accounting attributes used by the Service Manager application.

Table 158: Service Manager RADIUS Accounting Attributes

Attribute Number	Attribute Name	RADIUS Message Type	VSA Description
[26-83]	Service-Session	For service sessions only: Acct-Start Acct-Stop Interim-Acct	Name of the service (including parameter values) with which the statistics are associated
[26-140]	Service-Interim-Acct-Interval	Access-Accept and COA-Request	Number of seconds between accounting updates for a service; a tagged VSA

Table 158: Service Manager RADIUS Accounting Attributes (*continued*)

Attribute Number	Attribute Name	RADIUS Message Type	VSA Description
[44]	Acct-Session-ID	Acct-Start Acct-Stop Interim-Acct	Accounting identifier that makes it easy to match start and stop records in a log file; the format is extended to include a colon-separated value that uniquely identifies the subscriber session

Related Documentation

- [Service Session Profiles Overview on page 711](#)
- [Working with Service Session Profiles on page 712](#)
- [Understanding Service Manager RADIUS Attributes on page 686](#)
- [Service Interim Accounting Overview on page 699](#)
- [Configuring Service Interim Accounting on page 702](#)
- [Activating Subscriber Service Sessions Using RADIUS on page 685](#)
- [Deactivating Service Sessions Using RADIUS on page 690](#)

Service Interim Accounting Overview

Interim accounting determines how often accounting information is updated and sent to an accounting server. In addition to the user-based interim accounting supported on the router, Service Manager supports service-related interim accounting—you can configure an interim accounting interval for services that are created during a user RADIUS-based login and services that are activated by a COA operation.

The service interim accounting interval is specified by the RADIUS Service-Interim-Acct-Interval attribute (VSA 26-140) that is included in the RADIUS Access-Accept message or COA-Request message that activates a service session. Because the Service-Interim-Acct-Interval attribute is a tagged attribute, you can configure different interim accounting intervals for a particular user's various services.

You can use the **aaa service accounting interval** command to specify the default service interim accounting interval. Service Manager uses this interval value for service accounting when the Service-Interim-Acct-Interval attribute is not configured.



NOTE: You can also configure interim accounting for users. A user interim accounting interval is configured in the Acct-Interim-Interval RADIUS attribute (RADIUS attribute 85). You use the **aaa user accounting interval** command to specify the default user interim accounting interval, which is used when RADIUS attribute 85 is not configured. See the *Configuring Remote Access* chapter in this guide for information about configuring user interim accounting.

When the Service-Interim-Acct-Interval attribute is configured for a service, Service Manager uses the guidelines shown in [Table 159 on page 700](#) to determine the correct interim accounting interval to use for the service.

Table 159: Determining the Service Interim Accounting Interval

Service-Interim-Acct-Interval Value	Service Manager Action
0	Disables interim accounting for the service
1–599	Uses 600
600–86400	Uses the specified value
86401 or greater	Uses 86400
The tag for the service-interim-acct-interval attribute does not match the tag for any service-activate attribute (VSA 26-65)	Discards the service-interim-acct-interval attribute

[Table 160 on page 700](#) describes a sample Acct-Start message for a service session. In the table, the three fields used by Service Manager are shown in bold characters. An Acct-Start message for a subscriber session without any active services does not include the Service-Session attribute.

Table 160: Sample Acct-Start Message for a Service Session

RADIUS Attribute	Sample Value
acct-status-type	1
username	client1@isp1.com
event-timestamp	1112191723
acct-delay-time	0
nas-identifier	ERX-01-00-06
acct-session-id	erx FastEthernet 12/0:0001048580:002478
nas-ip-address	10.6.128.45
class	(binary data)
framed-protocol	0
framed-compression	0
framed-ip-address	100.20.0.1

Table 160: Sample Acct-Start Message for a Service Session (*continued*)

RADIUS Attribute	Sample Value
framed-ip-netmask	0.0.0.0
ingress-policy-name (vsa)	forwardAll
egress-policy-name (vsa)	forwardAll
calling-station-id	#ERX-01-00-06#E12#0
acct-input-gigawords	0
acct-input-octets	4032
acct-output-gigawords	0
acct-output-octets	2163
acct-input-gigapackets (vsa)	0
acct-input-packets	7
acct-output-gigapackets (vsa)	0
acct-output-packets	7
nas-port-type	15
nas-port	3221225472
nas-port-id	FastEthernet 12/0
acct-authentic	1
acct-session-time	0
acct-service-session	tiered(1280000, 5120000)
service-interim-acct-interval	1200

Related Documentation

- [Understanding Service Manager RADIUS Attributes on page 686](#)
- [Understanding RADIUS Accounting for Service Manager on page 697](#)
- [Service Interim Accounting for IPv4 and IPv6 Services in a Dual Stack Overview on page 705](#)
- [Configuring Service Interim Accounting on page 702](#)

Configuring Service Interim Accounting

You can configure user-based interim accounting on the router. In addition, you can configure service-related interim accounting for services that are created during a user RADIUS-based login and services that are activated by a COA operation.

Configuring service interim accounting comprises the following sets of tasks:

- [Specifying the Service Accounting Interval on page 702](#)
- [Specifying the User Accounting Interval on page 702](#)

Specifying the Service Accounting Interval

You can specify the default interval between service accounting updates by using the **aaa service accounting interval** command. Service manager uses the default interval when no value is specified in the Service-Interim-Acct-Interval attribute (Juniper VSA 26-140).

To specify the default interval between service accounting updates:

- Issue the **aaa service accounting interval** command in Global Configuration mode.

```
host1(config)#aaa service accounting interval 60
```

You can specify the service accounting interval in the range of 10–1440 minutes. The default setting is 0, which disables the feature.



NOTE: To enable interim service accounting, the service accounting interval must be set to a non-zero value and the service statistics type must *not* be set to *none*.

This command and the **aaa user accounting interval** command replace the **aaa accounting interval** command, which is deprecated and might be removed in a future release. The default interval is applied on a virtual router basis—this setting is used for services associated with all users who attach to the corresponding virtual router.

Use the **no** version to reset the accounting interval to 0, which turns off interim service accounting when no value is specified in the Service-Interim-Acct-Interval attribute (Juniper VSA 26-140).

Specifying the User Accounting Interval

You can specify the default interval between user accounting updates by using the **aaa user accounting interval** command. The router uses the default interval when no value is specified in the RADIUS Acct-Interim-Interval attribute (RADIUS attribute 85).

To specify the default interval between user accounting updates:

- Issue the **aaa user accounting interval** command in Global Configuration mode.

```
host1(config)#aaa user accounting interval 20
```

The default interval is applied on a virtual router basis—this setting is used for all users who attach to the corresponding virtual router. You can specify the user accounting interval, in the range 10–1440 minutes. The default setting is 0, which disables the feature.

This command and the **aaa service accounting interval** command replace the **aaa accounting interval** command, which is deprecated and might be removed in a future release.

Use the **no** version to reset the accounting interval to 0, which turns off interim user accounting when no value is specified in the RADIUS Acct-Interim-Interval attribute.

Related Documentation

- [Understanding Service Manager RADIUS Attributes on page 686](#)
- [Understanding RADIUS Accounting for Service Manager on page 697](#)
- [Service Interim Accounting Overview on page 699](#)
- [Service Interim Accounting for IPv4 and IPv6 Services in a Dual Stack Overview on page 705](#)
- *aaa service accounting interval*
- *aaa user accounting interval*

Configuring Calculation of Service Session Accounting Based on Scheduler Profiles Instead of Rate-Limit Profiles in Hierarchical Parent Groups for Forwarded Packets

You can configure the Service Manager application to compute accounting for service sessions based on the number of packets that are forwarded by the scheduler profile in QoS profiles on output interfaces. Service Manager collects statistics from parent groups that are configured in an output policy attached to an interface. Each parent group might contain a rate-limit profile. When packets enter an interface to which a rate-limit profile is applied, the router counts the number of bytes (packets) over time, categorizes each packet as committed, conformed, or exceeded, and assigns a transmit, drop, or mark action to the packets.

Parent groups configured in a hierarchy are effectively used in a layer 2 (ATM) access network for digital subscriber lines (DSLs) where many routing gateways lead to one Broadband Remote Access Server (B-RAS). The B-RAS uses rate-limit hierarchies to allocate shareable bandwidth to each routing gateway, which enables unused bandwidth from one routing gateway to be used by others.

Each rate-limit profile in a hierarchical parent group processes and classifies the packets that arrive at an interface, and computes the statistics to be displayed in the output of the **show** commands for the relevant interfaces. For rate-limit profiles in a hierarchical parent group that are part of classifier groups in a policy attached to an output interface, if you configure a scheduler profile in a QoS profile and attach it to the same output interface, the packets processed by the rate-limit profile in the hierarchical parent group might be either forwarded or dropped based on the scheduler profile attached to the interface. A scheduler profile configures the bandwidth at which a traffic queue is validated as a function of relative weight, assured rate, and shaping rate.

This method of operation causes discrepancies and inaccuracies in the accounting statistics for subscriber service sessions that the Service Manager application collects. Because packets that are treated as forwarded packets at the rate-limit profile in a hierarchical parent group might be classified in a different way at some other level in a rate-limit hierarchy, incorrect accounting statistics are retrieved for service sessions.

In an environment in which you apply policies, with rate-limit profiles defined in a hierarchical parent group, to interfaces that are also assigned with scheduler profiles, you can enable Service Manager to calculate accounting information for subscriber sessions based on the packets forwarded by scheduler profiles. You can use the **service-accounting-statistics scheduler-based** command in Global Configuration mode to enable this capability to compute accounting details based on scheduler profiles for policies with hierarchical rate-limit profiles on output interfaces. By default, this feature is not enabled. This functionality is effective only for packets that are forwarded at an output interface by a policy that contains a rate-limit profile in a hierarchical parent group and does not apply to packets that are dropped by the rate-limit profile.

When you enable or disable this feature, the statistical values collected for policies that contain rate-limit profiles in hierarchical parent groups on output interfaces are reset. This setting to calculate accounting details for service sessions based on scheduler profiles is preserved across an upgrade operation.

When you enable this functionality, the following operations are performed:

- The packet received at an output interface is processed by the rate-limit profile defined in a hierarchical parent group after classification. If the packet is marked to be transmitted based on the action specified in the rate-limit profile, the statistical counter to save the number of forwarded packets in the parent group is not incremented. When the packet arrives at the scheduler profile and is processed by it, if it is scheduled to be forwarded, the counter to store the number of forwarded packets for each rate-limit profile is incremented.
- If the incoming packet is determined to be dropped by the rate-limit profile in the parent group against which it is validated, the counter to record the number of dropped packets of the parent group is incremented.

If you enable scheduler profile–based computation of service session accounting, for IPv4, IPv6, or MPLS interfaces, the output of the **show ip interface**, **show ipv6 interface**, or **show mpls interface** commands display the forwarded packets and bytes fields, and dropped packets and bytes fields in the rate-limit-profile section under the IP, IPv6, or MPLS policy output headings for policies with hierarchical parent groups. The committed, conformed, exceeded, saturated, and unconditional packets and bytes fields are not displayed in the rate-limit-profile section in the output of these commands for policies with hierarchical parent groups.

To enable computation of accounting statistics for service sessions based on scheduler profiles for forwarded packets with rate-limit profiles in hierarchical parent groups on output interfaces:

- From Global Configuration mode, enable the capability to calculate accounting statistics for subscriber service sessions by using scheduler profiles instead of rate-limit profiles in hierarchical parent groups:

```
host1(config)#service-accounting-statistics scheduler-based
```

Use the **no** version of this command to disable the computation of accounting details based on scheduler profiles. In such a case, the accounting information is computed based on rate-limit profiles defined in hierarchical parent groups for policies on output interfaces.

Related Documentation

- [Verifying Computation of Service Session Accounting Based on Scheduler Profiles on page 773](#)
- *service-accounting-statistics scheduler-based*
- *show service-accounting-statistics*

Service Interim Accounting for IPv4 and IPv6 Services in a Dual Stack Overview

You can query the external parent group statistics similar to the statistics retrieved for classifier lists. You must specify the correct external parent group name and its corresponding hierarchical policy parameter for each of the input, output, and secondary-input statistics. You can specify a list of external parent groups along with hierarchical policy parameter for which statistics must be collected and sent to Service Manager for display in the Acct-Stop and Interim-Acct messages. The statistics for packets arriving at an interface attached at the input stage and the statistics for packets arriving at an interface attached at the secondary input stage are added and displayed in the Input Bytes field of the **show service-management** command. The statistics for packets leaving an interface at which the hierarchical policy is defined are displayed in the Output Bytes field of the **show service-management** command. The external parent group statistics are not limited to combined IPv4 and IPv6 services in a dual stack. You can also obtain external parent group statistics for IPv4 and IPv6 services configured independently in a dual stack.

You can retrieve either external parent group statistics or classifier statistics from policy manager. However, you cannot retrieve both the statistics for a single service definition. When a combined service is configured, you cannot retrieve classifier list-based based statistics. In such a scenario, you can only retrieve external parent group-based statistics from policy manager.

Service interim accounting and accounting based on service deactivation are supported for IPv6 services. For the combined IPv4 and IPv6 service, the statistics are a sum of the values in the external parent group and hierarchical policy parameter pair lists (defined as input-stat-epg, secondary-input-stat-epg, and output-stat-epg in the service definition macro).

If an interface fails, service-related interim accounting does not calculate the packets that are transmitted through this failed interface. For statistics reporting, only those packets that exist for interfaces when the subscriber service session is deactivated are counted.

Related Documentation

- [Understanding RADIUS Accounting for Service Manager on page 697](#)
- [Service Interim Accounting Overview on page 699](#)
- [Configuring Service Interim Accounting on page 702](#)
- [Combined and Independent IPv4 and IPv6 Services in a Dual Stack Overview on page 694](#)
- [Activation and Deactivation of IPv4 and IPv6 Services in a Dual Stack on page 696](#)
- [Example: Combined IPv4 and IPv6 Service in a Dual Stack Service Definition on page 740](#)

Collecting IPv6 Service Accounting Statistics Separately for Dual-Stack Subscribers Overview

In a dual-stack environment, you can send IPv6 service accounting statistics separately to the accounting server when either a combination of IPv4 and IPv6 subscribers or an IPv6 subscriber is configured. The Service Manager tracks the following JunosE objects passed back through the **env.setResult** method to collect IPv6 service accounting statistics:

- secondary-input-stat-clacl
- input-stat-clacl
- output-stat-clacl
- input-stat-epg
- output-stat-epg



.....

NOTE: Collecting the Service Manager statistics collection is a three-part procedure. You must configure the statistics information in the service definition macro file, enable the collection of the statistics by using either RADIUS or the command-line interface (CLI), and enable the collection of the statistics for the policy referenced in the service macro by using the **statistics enabled** keyword in the command for assigning the policy to a profile at the time of attaching the policy to an interface. For detailed information about Service Manager statistics, see [“Configuring Service Manager Statistics” on page 718](#).

.....

The collected statistics are sent to the accounting server separately through the Acct-Stop and Interim-Acct RADIUS messages by using IPv6 accounting RADIUS vendor-specific attributes (VSAs) if the inclusion of these attributes in RADIUS messages is enabled using the **radius include** command.



NOTE: For more information about the RADIUS VSAs, see [“Juniper Networks VSAs Supported for Subscriber AAA Accounting Messages” on page 196](#) and [“Juniper Networks VSAs” on page 248](#). For more information about the usage of the `radius include` command, see [“CLI Commands Used to Configure Juniper Networks VSAs” on page 213](#).

The following RADIUS VSAs are used for sending the IPv6 service accounting statistics to the accounting server:

- `Ipv6-Acct-Input-Octets` [26-151]
- `Ipv6-Acct-Output-Octets` [26-152]
- `Ipv6-Acct-Output-Octets` [26-152]
- `Ipv6-Acct-Output-Packets` [26-154]
- `Ipv6-Acct-Input-Gigawords` [26-155]
- `Ipv6-Acct-Output-Gigawords` [26-156]

**Related
Documentation**

- [Combined and Independent IPv4 and IPv6 Services in a Dual Stack Overview on page 694](#)
- [Activation and Deactivation of IPv4 and IPv6 Services in a Dual Stack on page 696](#)
- [Service Interim Accounting for IPv4 and IPv6 Services in a Dual Stack Overview on page 705](#)

Managing Subscriber Service Sessions by Using the CLI Overview

The CLI-based Service Manager creates static subscriber sessions and service sessions. You can also use CLI commands to immediately deactivate subscriber service sessions. The CLI-based support is particularly useful for:

- Testing your service definitions—for example, you might use the CLI commands to verify that a newly created service definition is correct. When you are satisfied with the service definition, you can then use RADIUS to activate the service for your subscribers.
- Preprovisioning Service Manager services—preprovisioning improves performance and efficiency by freeing Service Manager from having to repeatedly create and remove a service that you activate and deactivate for multiple subscribers. See [“Preprovisioning Service Sessions” on page 711](#) for more information about service preprovisioning.

**Related
Documentation**

- [Service Session Profiles Overview on page 711](#)
- [Working with Service Session Profiles on page 712](#)
- [Managing and Activating Service Sessions Overview on page 684](#)
- [Activating Subscriber Service Sessions by Using the CLI Overview on page 708](#)
- [Activating Subscriber Sessions by Using the CLI on page 709](#)

- [Preprovisioning Service Sessions on page 711](#)
- [Deactivating Subscriber Service Sessions by Using the CLI Overview on page 715](#)

Activating Subscriber Service Sessions by Using the CLI Overview

A subscriber session represents a specific subscriber—the session consists of the subscriber’s name, the interface used for the session, and any active services for the subscriber. A subscriber can have one subscriber session active at any given time.

You create a subscriber’s service session when you assign a service definition to a subscriber session. Like an AAA-created service, a single subscriber session can have multiple simultaneous service sessions. You can use one method to create the subscriber session, and then a different method to activate the subscriber’s service session. For example, you might use RADIUS to create the AAA subscriber session, then use the CLI to activate the service session for the subscriber. You can optionally specify a service session profile that you want to attach to the service session.

You can use the CLI to activate a service session based on subscriber information or owner information:

- Subscriber name and interface method—Activates the service session based on the subscriber name and the interface that the subscriber is using for this subscriber session.
- Owner name and ID method—Activates the service session based on the owner that created the subscriber session and the ID that was generated by the owner. For example, if RADIUS is used to create the subscriber session, the owner name is AAA and the owner ID is the Acct-Session-ID that was generated by RADIUS during subscriber creation.



NOTE: You must specify the parameter values in the order in which the parameters appear in the template name of the service definition file. Enclose the service definition name in double quotation marks, with the service’s parameter values in parentheses. For example, for the tiered service that is defined in “[Creating Service Definitions](#)” on page 668, the template name is:

`<# tiered(inputBW, outputBW) #>`

Use the following format with the **service-session** keyword:

`“tiered(1280000, 5120000)”`

Related Documentation

- [Service Session Profiles Overview on page 711](#)
- [Managing and Activating Service Sessions Overview on page 684](#)
- [Managing Subscriber Service Sessions by Using the CLI Overview on page 707](#)
- [Activating Subscriber Sessions by Using the CLI on page 709](#)

- [Deactivating Subscriber Service Sessions by Using the CLI Overview on page 715](#)
- [Gracefully Deactivating Subscriber Service Sessions on page 716](#)
- [Forcing Immediate Deactivation of Subscriber Service Sessions on page 717](#)

Activating Subscriber Sessions by Using the CLI

You can use the CLI to activate a service session based on subscriber information or owner information. You can activate a service session by performing one of the following tasks:

- [Activating a Service for an Existing Subscriber on page 709](#)
- [Creating and Activating a Service for a Subscriber on page 710](#)

Activating a Service for an Existing Subscriber

You can use the **service-management owner-session** command to activate a service for an existing subscriber by identifying the owner used to create the subscriber session and specifying the service session to use. The subscriber session must exist before you use this command.

You can use this command in Privileged Exec mode to create a dynamic subscriber session—dynamic sessions are deleted after a router reboot. You can use this command in Global Configuration mode to create persistent subscriber sessions that are retained across reboots.

You can specify the name of the owner (the method originally used to create the subscriber session), and the ID generated by the owner. For example, if RADIUS was used to create the subscriber session, the owner name is AAA and the owner ID is the Acct-Session-ID generated by RADIUS when the subscriber session was created.

You can activate one subscriber session for a subscriber—and multiple service sessions for a particular subscriber session. If you create a second subscriber session for the same subscriber, only the newest subscriber session, with its services, is used.

You can include the optional **service-session-profile** keyword to assign a profile to the service session. The service session profile includes additional attributes, such as the type of statistics to be captured for the service session.

- To activate a service session for an existing subscriber:

```
host1(config)#service-management owner-session aaa 573498 service-session
"video(4500000,192.168.10.3)"
```

- To activate multiple service sessions for an existing subscriber:

```
host1(config)#service-management owner-session aaa 573498 service-session
"video(4500000,192.168.10.3)"
host1(config)#service-management owner-session aaa 573498 service-session
"tiered(1000000,2000000)"
host1(config)#service-management owner-session aaa 573498 service-session
"voice(1000000,10.10.10.1)"
```

- To include a service session profile when you activate a subscriber's service session:

```
host1(config)#service-management owner-session aaa 426777 service-session
"video(4500000, 192.168.10.3)" service-session-profile vodISP1
```

Use the **no** version to deactivate service sessions based on owner information. See ["Deactivating Subscriber Service Sessions by Using the CLI Overview"](#) on page 715 for more information about deactivating subscriber service sessions.

Creating and Activating a Service for a Subscriber

You can use the **service-management subscriber-session service-session** command to activate a service for a subscriber by creating a subscriber session and a service session.



NOTE: Always activate at least one service session for a subscriber session. The ability to create a subscriber session without a service session (by omitting the **service-session** keyword) is not currently supported.

You can use this command in Privileged Exec mode to create a dynamic subscriber session—dynamic sessions are deleted after a router reboot. You can use this command in Global Configuration mode to create persistent subscriber sessions that are retained across reboots.

You can create one subscriber session for a subscriber—and multiple service sessions for a particular subscriber session. If you create a second subscriber session for the same subscriber, only the newest subscriber session, with its services, is used.

You can include the optional **service-session-profile** keyword to assign a profile to the service session. The service session profile includes additional attributes, such as the type of statistics to be captured for the service session.

- To activate a subscriber session with a single service session:

```
host1(config)#service-management subscriber-session client1@isp1.com interface
atm 4/0.1 service-session "video(4500000, 192.168.10.3)"
```

- To activate a single subscriber session with multiple service sessions:

```
host1(config)#service-management subscriber-session client1@isp1.com interface
atm 4/0.1 service-session "video(4500000, 192.168.10.3)"
host1(config)#service-management subscriber-session client1@isp1.com interface
atm 4/0.1 service-session "tiered(1000000, 2000000)"
host1(config)#service-management subscriber-session client1@isp1.com interface
atm 4/0.1 service-session "voice(1000000, 10.10.10.1)"
```

- To include a service session profile when you activate a subscriber's service session:

```
host1(config)#service-management subscriber-session client1@isp1.com interface
atm 4/0.1 service-session "video(4500000, 192.168.10.3)" service-session-profile
vodISP1
```

- Use the **no** version to deactivate service sessions. See ["Deactivating Subscriber Service Sessions by Using the CLI Overview"](#) on page 715 for more information about deactivating subscriber service sessions.

- Related Documentation**
- [Service Session Profiles Overview on page 711](#)
 - [Working with Service Session Profiles on page 712](#)
 - [Managing Subscriber Service Sessions by Using the CLI Overview on page 707](#)
 - [Activating Subscriber Service Sessions by Using the CLI Overview on page 708](#)
 - [Preprovisioning Service Sessions on page 711](#)
 - *service-management subscriber-session service-session*

Preprovisioning Service Sessions

Preprovisioning service sessions is a technique you can use to improve Service Manager's performance. Typically, when you use a service definition to activate a subscriber's service session, Service Manager uses resources to build that service. However, if you later use the same service definition to activate a service session for a second subscriber, Service Manager does not have to rebuild the service—it bases the new service on the service that it built for the first service session. After you deactivate the first session, Service Manager must build a new service for any subsequent subscribers.

Preprovisioning entails activating a service for a dummy user on the null interface. You can then use the preprovisioned service session to activate service sessions for actual subscribers. This technique improves performance because you only require Service Manager to build the service one time, then reuse the original service when you activate future subscriber service sessions.

To preprovision a service you use a command similar to the following example:

```
host1(config)#service-management subscriber-session dummy interface null
service-session "tiered(1000000,2000000)"
```

- Related Documentation**
- [Service Session Profiles Overview on page 711](#)
 - [Working with Service Session Profiles on page 712](#)
 - [Activating Subscriber Sessions by Using the CLI on page 709](#)
 - [Deactivating Subscriber Service Sessions by Using the CLI Overview on page 715](#)
 - [Gracefully Deactivating Subscriber Service Sessions on page 716](#)
 - [Forcing Immediate Deactivation of Subscriber Service Sessions on page 717](#)
 - [Using Service Session Profiles to Deactivate Service Sessions on page 717](#)

Service Session Profiles Overview

Service session profiles provide additional flexibility to the Service Manager application by enabling you to assign one or more supported attributes to a particular activation of a service.

For example, you might assign the same video service to two subscribers, but use different service session profiles to set different time limits for each subscriber's service. One subscriber uses the video service for 5 hours (18000 seconds) while the other subscriber's video service is for 10 hours (36000 seconds). Or, you might enable statistics on a subscriber's voice service and disable statistics on the same subscriber's video service.

You can create multiple service session profiles independent of the service activation process. Then, when you activate a service session, you specify the profile that you want to use with that particular service session—you can apply one profile to a service session.

You can configure the following attributes in service session profiles:

- **statistics**—Enables statistics and specifies the type of statistics you want to capture for the service. See [“Configuring Service Manager Statistics” on page 718](#) for additional information about capturing Service Manager statistics. You can specify the following types of statistics:
 - **time**—The service's duration
 - **volume-time**—The service's duration and traffic volume
- **volume**—Specifies that the service is automatically deactivated when the indicated traffic volume is exceeded.
- **time**—Specifies that the service is automatically deactivated when the indicated time period is exceeded.



NOTE: The **volume** and **time** attributes use values captured by the Service Manager statistics feature to determine when the threshold is exceeded. Service Manager collects time statistics by default—you must configure and enable volume statistics collection. See [“Configuring Service Manager Statistics” on page 718](#).

Related Documentation

- [Working with Service Session Profiles on page 712](#)
- [Managing and Activating Service Sessions Overview on page 684](#)
- [Managing Subscriber Service Sessions by Using the CLI Overview on page 707](#)
- [Activating Subscriber Service Sessions by Using the CLI Overview on page 708](#)

Working with Service Session Profiles

This topic describes the settings you can configure for a new service session profile or the attributes you can modify in an existing service session profile. Creating or modifying a service session profile comprises the following sets of tasks:



NOTE: To modify an existing profile, you can add new attributes or use the **no** version of a command to remove an attribute.

- [Creating a New Service Session Profile on page 713](#)
- [Specifying Statistics Collection Settings on page 713](#)
- [Specifying the Maximum Bandwidth for a Service Session on page 714](#)
- [Specifying the Interval for the Active State of a Service Session on page 714](#)

Creating a New Service Session Profile

You can create a new service session profile, and enter Service Session Profile Configuration mode by using the **service-management service-session-profile** command. You can also use this command to modify the attributes of a previously created profile.

To create a new service session profile:

- Issue the **service-management service-session-profile** command in Global Configuration mode. Issuing this command enables you to access Service Session Profile Configuration mode.

```
host1(config)#service-management service-session-profile vodISP1
host1(config-service-session-profile)#
```

In Service Session Profile Configuration mode, you specify the attributes used in the service session profile, such as the maximum volume limit for the session and the maximum time the session can be used. You can also specify that Service Manager collect statistics for time, or volume, or both.

Use the **no** version to delete the service session profile.

Specifying Statistics Collection Settings

You can use the **statistics** command in Service Session Profile Configuration mode to enable statistics collection and to specify the type of statistics to collect.

- Use the **time** keyword with the **statistics** command to collect statistics about the duration of the service session.

```
host1(config)#service-management service-session-profile vodISP1
host1(config-service-session-profile)#statistics time
```

- Use the **volume-time** keyword with the **statistics** command to collect statistics about both the volume of traffic and the duration of the service session.

```
host1(config)#service-management service-session-profile vodISP1
host1(config-service-session-profile)#statistics volume-time
```

Use the **no** version to disable statistics collection.



NOTE: Service Manager statistics collection is a three-part procedure. You must configure statistics information in the service definition macro file, enable statistics collection by either RADIUS or the CLI, and also enable statistics collection for the policy referenced in the service macro using the **statistics enabled** keyword in the command used for policy attachment in the profile. See [“Configuring Service Manager Statistics” on page 718](#).

Specifying the Maximum Bandwidth for a Service Session

You can use the **volume** command in Service Session Profile configuration mode to specify the maximum amount of bandwidth that can use the service.

To configure the maximum amount of bandwidth for a service session:

- Issue the **volume** command in Service Session Profile configuration mode.

```
host1(config)#service-management service-session-profile vodISP1
host1(config-service-session-profile)#volume 1000000
```

The router immediately terminates the subscriber's service session when the specified traffic volume is exceeded. The range is 0–16777251MB.

Use the **no** version to delete the volume attribute from the service session profile.



NOTE: The **volume** attribute uses values captured by the Service Manager statistics feature to determine when the threshold is exceeded. Therefore, you must configure and enable statistics collection to use this attribute. See [“Configuring Service Manager Statistics” on page 718](#).

Specifying the Interval for the Active State of a Service Session

You can use the **time** command in Service Session Profile configuration mode to specify the maximum amount of time that the service session can be active for the subscriber.

To configure the maximum duration for which a subscriber service session can be active:

- Issue the **time** command in Service Session Profile configuration mode.

```
host1(config)#service-management service-session-profile vodISP1
host1(config-service-session-profile)#time 6000
```

The router immediately terminates the subscriber's service session when the specified time is exceeded. The range is 0–16777251 seconds.

Use the **no** version to delete the time attribute from the service session profile.

Related Documentation

- [Service Session Profiles Overview on page 711](#)
- [Activating Subscriber Sessions by Using the CLI on page 709](#)

- [Deactivating Subscriber Service Sessions by Using the CLI Overview on page 715](#)
- [Preprovisioning Service Sessions on page 711](#)
- [Gracefully Deactivating Subscriber Service Sessions on page 716](#)
- [Forcing Immediate Deactivation of Subscriber Service Sessions on page 717](#)
- [Using Service Session Profiles to Deactivate Service Sessions on page 717](#)
- *service-management service-session-profile*
- *statistics*
- *time*
- *volume*

Deactivating Subscriber Service Sessions by Using the CLI Overview

The CLI supports several methods that enable you to manually deactivate service sessions. You can:

- Gracefully terminate all services or a specific service for a particular subscriber
- Gracefully terminate all service or a specific service associated with a particular owner
- Force the immediate termination of all of a subscriber's sessions
- Use service session profiles to create time or volume thresholds for the service and deactivate the service when the threshold is reached. See [“Service Session Profiles Overview” on page 711](#) and [“Working with Service Session Profiles” on page 712](#).



NOTE: You can use the CLI commands described in this topic to delete subscriber and service sessions that are created by either RADIUS or the CLI.

The Service Manager CLI commands enable you to use variations of the **no service-management subscriber-session** command to terminate service sessions.

Related Documentation

- [Service Session Profiles Overview on page 711](#)
- [Working with Service Session Profiles on page 712](#)
- [Gracefully Deactivating Subscriber Service Sessions on page 716](#)
- [Forcing Immediate Deactivation of Subscriber Service Sessions on page 717](#)
- [Using Service Session Profiles to Deactivate Service Sessions on page 717](#)

Gracefully Deactivating Subscriber Service Sessions

You can deactivate a specific service for a subscriber, or you can delete a subscriber session, which deactivates all of the subscriber's service sessions. We recommend you use this command to deactivate subscriber service sessions.

You can gracefully deactivate subscriber service sessions using one of the following tasks:

- [Gracefully Deactivating Service Sessions Based on Owner Details on page 716](#)
- [Gracefully Deactivating Service Sessions Based on Subscriber Details on page 716](#)

Gracefully Deactivating Service Sessions Based on Owner Details

You can use the **no service-management owner-session** command to gracefully deactivate service sessions for a subscriber based on owner information.

To gracefully deactivate service sessions for a subscriber based on owner information:

- Issue the **no service-management owner-session** command in Global Configuration mode by specifying the owner name and owner ID of the service session you want to deactivate.

```
host1(config)#no service-management owner-session aaa 426777 service-session  
"video(4500000,192.168.10.3)"
```

You can use the **no** version with the **service-session** keyword to deactivate the specified service session. You can use the **no** version *without* the **service-session** keyword to delete the subscriber's session and deactivate all of the subscriber's service sessions.

Gracefully Deactivating Service Sessions Based on Subscriber Details

You can use the **no service-management subscriber-session** command to gracefully deactivate service sessions for a subscriber based on the subscriber credentials.

To gracefully deactivate service sessions for a subscriber based on the subscriber's username and interface, and not the subscriber session ID:

- Issue the **no service-management subscriber-session** command in Global Configuration mode by specifying the owner name and owner ID of the service session you want to deactivate.

```
host1(config)#no service-management subscriber-session client1@isp1.com interface  
atm 4/0.1 service-session "tiered(1000000,2000000)"
```

You can use the **no** version without the **service-session** keyword to delete the subscriber's session and deactivate all of the subscriber's service sessions. You can use the **no** version with the **service-session** keyword to deactivate the specified service session.

Related Documentation

- [Service Session Profiles Overview on page 711](#)
- [Working with Service Session Profiles on page 712](#)

- [Activating Subscriber Sessions by Using the CLI on page 709](#)
- [Deactivating Subscriber Service Sessions by Using the CLI Overview on page 715](#)
- [Forcing Immediate Deactivation of Subscriber Service Sessions on page 717](#)
- [Using Service Session Profiles to Deactivate Service Sessions on page 717](#)

Forcing Immediate Deactivation of Subscriber Service Sessions

You can force the immediate deactivation of the specified subscriber session. Such an action deletes all active service sessions for the subscriber. We recommend this method if you encounter difficulty when you used the graceful deactivation method. Always use the graceful method first.

To force the immediate termination of a subscriber session and to deactivate all services for the specified subscriber session:

- Issue the **no service-management subscriber-session** command in Global Configuration mode by specifying the subscriber session ID with the **force** keyword.

```
host1(config)#no service-management subscriber-session 8 force
```



NOTE: To determine the subscriber session ID of a session you want to deactivate, use the **show service-management subscriber-session brief** command. The display lists the IDs of all active subscriber sessions and the owner that created the session, such as AAA (RADIUS) or CLI.

There is no affirmative version of this command; there is only a **no** version.

Related Documentation

- [Service Session Profiles Overview on page 711](#)
- [Working with Service Session Profiles on page 712](#)
- [Activating Subscriber Sessions by Using the CLI on page 709](#)
- [Preprovisioning Service Sessions on page 711](#)
- [Deactivating Subscriber Service Sessions by Using the CLI Overview on page 715](#)
- [Gracefully Deactivating Subscriber Service Sessions on page 716](#)
- [Using Service Session Profiles to Deactivate Service Sessions on page 717](#)
- *service-management owner-session*
- *service-management subscriber-session service-session*

Using Service Session Profiles to Deactivate Service Sessions

To terminate a subscriber service session when a threshold is reached, you create a service session profile that includes a time threshold, or a volume threshold, or both.

Then, you attach the service session profile when you activate the service session. When the specified threshold is reached, the service session is terminated.



NOTE: This feature is not supported by the **service-management owner-session** command. The **service-management owner-session** command only supports service session profiles when activating service sessions.

The following procedure shows the commands you might use to create a time threshold for deactivating a service session. See [“Working with Service Session Profiles” on page 712](#) for information about using the **time** and **volume** keywords in service session profiles.

To create or modify a service session profile:

1. Specify the name of the service session profile and configure the threshold:

```
host1(config)#service-management service-session-profile vodISP1
host1(config-service-session-profile)#time 6000
host1(config-service-session-profile)#exit
```

2. Include the service session profile when you activate the subscriber service session:

```
host1(config)#service-management subscriber-session client1@isp1.com interface
atm 4/0.1 service-session “video(4500000, 192.168.10.3)” service-session-profile
vodISP1
```

Related Documentation

- [Service Session Profiles Overview on page 711](#)
- [Working with Service Session Profiles on page 712](#)
- [Activating Subscriber Sessions by Using the CLI on page 709](#)
- [Preprovisioning Service Sessions on page 711](#)
- [Deactivating Subscriber Service Sessions by Using the CLI Overview on page 715](#)
- [Gracefully Deactivating Subscriber Service Sessions on page 716](#)
- [Forcing Immediate Deactivation of Subscriber Service Sessions on page 717](#)
- *service-management owner-session*
- *service-management subscriber-session service-session*

Configuring Service Manager Statistics

The Service Manager application provides a flexible and efficient process for identifying and capturing statistics related to subscriber service sessions. Configuring Service Manager to collect statistics is a three- part process. First, you design the service definition macro file to identify the statistics that you want to collect. Second, you configure Service Manager to enable statistics collection when a service session is activated by either RADIUS or the CLI. Third, before you reference a policy in the service definition macro to enable Service Manager collect statistics, you must enable statistics collection for this

policy using the **statistics enabled** keyword in the command used for policy attachment in the profile.

The following topics describe how to configure the service definition macro file. For information about configuring Service Manager to enable statistics, see the *Enabling Statistics Collection with the CLI* section if you are using RADIUS to activate services, or see *Enabling Statistics Collection with RADIUS* section if you are using the CLI.

Configuring Service Manager statistics involves the following sets of tasks:

- [Setting Up the Service Definition File for Statistics Collection on page 719](#)
- [Enabling Statistics Collection with RADIUS on page 720](#)
- [Enabling Statistics Collection with the CLI on page 721](#)
- [Setting Up the External Parent Group Statistics Collection on page 722](#)

Setting Up the Service Definition File for Statistics Collection

Service Manager statistics are based on classifier lists—the classifier lists are referenced by policy lists that you define in your service definition macro file.

When you configure your service definition for statistics, you include the macro environment command **env.setResult** to indicate the type of statistics to track and to identify the classifier lists to use when generating statistics. The format of the environment command is:

```
<# env.setResult("string", "classifier-list-name" ) #>
```

The *string* variable specifies the type of statistics to track—Service Manager supports the following strings:

- **input-stat-clacl**—track input statistics
- **output-stat-clacl**—track output statistics
- **secondary-input-stat-clacl**—track input statistics for a policy attached at the secondary input stage

You can use the **input-stat-clacl**, **output-stat-clacl**, and **secondary-input-stat-clacl** statistics types to collect both IPv4 and IPv6 service accounting statistics independently.

The *classifier-list-name* variable is the name of the classifier list that is associated with the policy list that is defined in the service definition. You can specify multiple classifier lists in the command.

This example is a portion of the service definition macro file in [“Creating Service Definitions” on page 668](#). The two highlighted commands specify the statistics used by the Service Manager application.

```
profile <# name; '\n' #>
  ip policy secondary-input <# name #> statistics enabled merge
  ip policy output <# oname #> statistics enabled merge
  qos-profile triplePlayIP
  qos-parameter maxSubscBW <# outputBW; '\n' #>
```

```
<# env.setResult("activate-profile", name) #> <#
env.setResult("secondary-input-stat-clacl", "matchAll") #> <#
env.setResult("output-stat-clacl", "matchAll") #>
<# endtmpl #>
```

The `<# env.setResult("secondary-input-stat-clacl", "matchAll") #>` command specifies that Service Manager track statistics associated with the classifier list named `matchAll`, and that this classifier list is associated with the policy that is attached at the secondary input stage.

The `<# env.setResult("output-stat-clacl", "matchAll") #>` command specifies that Service Manager track the output statistics associated with the `matchAll` classifier list, which is associated with the policy attached at the output stage.



NOTE: Before you reference as policy in the service definition macro to enable Service Manager collect statistics, you must enable statistics collection for this policy using the `statistics enabled` keyword in the command used for policy attachment in the profile.

This example shows how you can also configure your service definition to collect total statistics from multiple classifier lists. The following command specifies that three classifier lists are used to generate output statistics for a service created by the service definition. Each time statistics are reported for this service, Service Manager uses the total of the statistics for `clacl1`, `clacl2`, and `clacl3`.



NOTE: The multiple classifier list can contain a combination of IPv4 and IPv6 classifiers.

```
<# env.setResult("output-stat-clacl", "clacl1 clacl2 clacl3" ) #>
```

Enabling Statistics Collection with RADIUS

You use the Service-Statistics RADIUS VSA [26-69] with either the RADIUS login or COA-Request method to enable statistics for RADIUS-activated service sessions. To enable statistics, configure the Service-Statistics VSA with a value of either 1 (timestamp only) or 2 (volume and timestamp).

[Table 161 on page 720](#) describes a partial RADIUS message in which the Service-Statistics attribute has a value of 2—this enables volume and timestamp statistics for the tiered service assigned to subscriber `client1@isp1.com`.

Table 161: RADIUS-Enabled Statistics

RADIUS Attribute	Tag	Value
username	none	client1@isp1.com
activate-service	6	tiered(1280000, 5120000)

Table 161: RADIUS-Enabled Statistics (*continued*)

RADIUS Attribute	Tag	Value
service-statistics	6	2

When you enable statistics for a RADIUS-activated service, RADIUS accounting reports can use the statistics.

Enabling Statistics Collection with the CLI

You use service session profiles to enable statistics when you activate a service session with the CLI. See [“Service Session Profiles Overview” on page 711](#) for detailed information about creating and using service session profiles.

For example, you can use the following procedure to capture statistics that are defined in the service definition macro file for the tiered service:

1. Configure the service session profile to enable statistics. Specify the type of statistics you want to capture (either time or both volume and time).

```
host1(config)#service-management service-session-profile isp1_tiered3
host1(config-service-session-profile)#statistics volume-time
host1(config-service-session-profile)#
```

2. Apply the service session when you activate the subscriber service session.

```
host1(config)#service-management subscriber-session client1@isp1.com interface
atm 4/0.1 service-session “tiered(1000000,2000000)” service-session-profile
isp1_tiered3
```

The captured statistics are now used when you use the Service Manager **show service-management** commands. For example:

```
host1# show service-management subscriber-session client1@isp1.com interface atm 4/0.1
service-session
User Name: client1@isp1.com, Interface: atm 4/0.1
Service : tiered(1000000,2000000)
Non-volatile : False
Owner : CLI
State : Config ApplySuccess
Activate : True
Statistics Type : time-based and volume-based
Statistics Complete : False
Poll Interval : 0
Poll Expire : 0
Activate Time : THU MAR 01 21:09:12 2006
Time : 0
Time Expire : 0
Volume MB: 2
Volume Expire MB: 1
Input Bytes : 594
Output Bytes : 1196
Input Packets : 1
Output Packets : 2
```

Setting Up the External Parent Group Statistics Collection

Policies for interface groups include external parent groups that are implicitly instantiated during policy attachment based on each unique interface group encountered. You can use external parent groups and policy parameters for sharing aggregate nodes across policy attachments. Each external parent group reference in a policy is accompanied by a parameter that is resolved during the attachment of the policy to an interface.

You can retrieve either external parent group statistics or classifier statistics from policy manager. However, you cannot retrieve both statistics for a single service definition. When a combined service is configured, you cannot retrieve classifier list-based statistics. In such a scenario, you can only retrieve external parent group-based statistics from policy manager.

Only hierarchical policy parameters can have external parent group references. Each parameter has a single value, depending on the type of parameter. The hierarchical policy parameter can have a single numeric value or a keyword. When you configure your service definition for statistics, you include the macro environment command `env.setResult` to indicate the type of statistics to track and to identify the external parent groups to use when generating statistics. The format of the environment command is:

```
<# env.setResult("string" , "external-parent-grp-name policy-parameter-name") #>
```

The *string* variable specifies the type of statistics to track. Service Manager supports the following strings:

- **input-stat-epg**—Track input statistics for an external parent group in a hierarchical policy
- **output-stat-epg**—Track output statistics for an external parent group in a hierarchical policy
- **secondary-input-stat-epg**—Track input statistics for an external parent group in a hierarchical policy attached at the secondary input stage

The *external-parent-grp-name* variable is the name of the external parent group that a classifier group refers to in the policy list that is defined in the service definition. You must specify the external parent group and the hierarchical policy in the `env.setResult` command as a pair. You can specify multiple pairs of external parent groups and hierarchical policies in the command. The *policy-parameter-name* variable is the name of the hierarchical policy that allows classifier groups and parent groups within a policy to point to line module global parent groups. Each reference to a policy parameter in a policy is substituted with its value for all attachments of this policy at the interface.

For example, if *v4v6* is the name of the hierarchical policy parameter and the external parent group names are *vc-v4v6-in* and *vc-v4v6-out*, you must configure both the external parent group names and the corresponding hierarchical policy parameter in the `env.setResult` method for the external parent group statistics to be calculated.

```
<# env.setResult("input-stat-epg", "vc-v4v6-in v4v6" ) #>  
<# env.setResult("output-stat-epg", "vc-v4v6-out v4v6" ) #>
```

The `<# env.setResult("secondary-input-stat-epg", "vc-v4v6-in v4v6") #>` command specifies that Service Manager track statistics associated with the external parent group named `vc-v4v6-in` and the corresponding hierarchical policy named `v4v6`, and that this external parent group is associated with the policy that is attached at the input stage.

The `<# env.setResult("output-stat-epg", "vc-v4v6-out v4v6") #>` command specifies that Service Manager track the output statistics associated with the external parent group named `vc-v4v6-out` and the corresponding hierarchical policy named `v4v6`, which is associated with the policy attached at the output stage.

The input and output statistics associated with the external parent group are collected and forwarded to the Service Manager to be displayed in the `Acct-Stop` and `Interim-Acct` messages.

Related Documentation

- [Service Session Profiles Overview on page 711](#)
- [Working with Service Session Profiles on page 712](#)
- [Service Definitions Overview on page 667](#)
- [Understanding Service Manager RADIUS Attributes on page 686](#)
- [Understanding RADIUS Accounting for Service Manager on page 697](#)
- *service-management service-session-profile*
- *service-management subscriber-session service-session*
- *show service-management service-session-profile*

Service Manager Performance Considerations

Like any application, Service Manager requires a certain amount of system resources. Consider the following guidelines to maximize the performance of Service Manager when delivering subscriber services:

- **Minimize service definitions**—Use the minimum number of JunosE commands in a service definition to specify a service.
- **Reference objects in service definitions**—Referencing commonly used objects is more resource-efficient than using unique objects for each subscriber (for example, using a subscriber's IP address as a match criteria in a classifier list).
- **Preprovision frequently used services**—Preprovisioning saves resources by requiring Service Manager to build a popular service only once. You then reuse the original service when you activate future subscriber service sessions. See [“Preprovisioning Service Sessions” on page 711](#) for details.
- **Capture volume statistics when needed**—Repeatedly capturing volume statistics can waste resources.

Related Documentation

- [Service Definitions Overview on page 667](#)
- [Referencing Policies in Service Definitions Overview on page 673](#)

- [Creating Service Definitions on page 668](#)
- [Service Session Profiles Overview on page 711](#)
- [Provisioning Service Sessions on page 711](#)

Example: Tiered Service Definition

The following example shows you how to create a tiered service.

- [Requirements on page 724](#)
- [Overview on page 724](#)
- [Creating a Tiered Service on page 725](#)

Requirements

This example uses the following software and hardware components:

- JunosE Release 7.1.0 or higher-numbered releases
- E Series router (ERX7xx models, ERX14xx models, the ERX310 router, the E120 router, or the E320 router)
- ASIC-based line modules that support Fast Ethernet or Gigabit Ethernet

Also, this example assumes that QoS profile triplePlayIP and QoS parameter maxSubscBW are configured.

Overview

A tiered service typically provides set bandwidths for both inbound and outbound traffic for a subscriber. In this example, the bandwidth values are parameterized.

Creating a Tiered Service

Step-by-Step Procedure This example illustrates how to create a tiered service:



NOTE: Commented text explains the parameterized values in the example of the service definition macro file. Each example is followed by examples of RADIUS information and the CLI command that you can use to activate a subscriber service session.

```
!parameterizes input and output bandwidth
<# tiered(inputBW, outputBW) #>

<# uid := app.servicemanager.getUniqueId #>
<# name := "SM-tiered-" $ uid #>
<# oname := "SM-O-tiered-" $ uid #>
classifier-list matchAll ip any any
rate-limit-profile <# name #> one-rate
    committed-rate <# inputBW; '\n' #>

policy-list <# name; '\n' #>
    classifier-group matchAll precedence 10000
    rate-limit-profile <# name; '\n' #>
    traffic-class best-effort

policy-list <# oname; '\n' #>
    classifier-group matchAll precedence 10000
    traffic-class best-effort

profile <# name; '\n' #>
    ip policy secondary-input <# name #> statistics enabled merge
    ip policy output <# oname #> statistics enabled merge
    qos-profile triplePlayIP
    qos-parameter maxSubscBW <# outputBW; '\n' #>

<# env.setResult("activate-profile", name) #>
<# env.setResult("secondary-input-stat-clacl", "matchAll") #>
<# env.setResult("output-stat-clacl", "matchAll") #>

<# endtmp1 #>
```

Sample RADIUS Attributes

Step-by-Step Procedure Table 162 on page 725 provides the sample RADIUS attributes that you can use to activate the tiered service

Table 162: Sample RADIUS Attributes

RADIUS Attribute	Tag	Value
username	none	client1@isp1.com
activate-service	1	tiered(1280000, 5120000)

Sample CLI Command

Step-by-Step Procedure

To activate a subscriber session with the configured tiered service:

```
host1(config)#service-management subscriber-session client1@isp1.com interface atm 4/0.1 service-session "tiered(1280000, 5120000)"
```

Related Documentation

- [Service Definitions Overview on page 667](#)
- [Creating Service Definitions on page 668](#)
- [Managing Your Service Definitions on page 671](#)
- [Understanding Service Manager RADIUS Attributes on page 686](#)
- [Understanding RADIUS Accounting for Service Manager on page 697](#)
- *service-management subscriber-session service-session*

Example: Video-on-Demand Service Definition

The following example shows a sample service definition macro file that creates a video-on-demand service—the service provides bandwidth that meets the needs of video streams.

- [Requirements on page 726](#)
- [Overview on page 726](#)
- [Creating a Video-on-Demand Service on page 727](#)

Requirements

This example uses the following software and hardware components:

- JunosE Release 7.1.0 or higher-numbered releases
- E Series router (ERX7xx models, ERX14xx models, the ERX310 router, the E120 router, or the E320 router)
- ASIC-based line modules that support Fast Ethernet or Gigabit Ethernet

Overview

The video-on-demand service definition creates the bandwidth towards the subscriber and parameterizes the source of the video feed.

Creating a Video-on-Demand Service

Step-by-Step Procedure This example illustrates how to create a video-on-demand service:



NOTE: Commented text explains the parameterized values in the example of the service definition macro file. Each example is followed by examples of RADIUS information and the CLI command that you can use to activate a subscriber service session.

```
!parameterizes download bandwidth and server address
<# videoMin(downloadBW, serverAddress) #>

<# uid := app.servicemanager.getUniqueId #>
<# name := "SM-video-" $ uid #>

classifier-list <# name #> ip any <# serverAddress #> 0.0.0.0

policy-list <# name; '\n' #>
  classifier-group <# name #> precedence 5000
  traffic-class video

profile <# name; '\n' #>
  ip policy output <# name #> statistics enabled merge
  qos-parameter maxVideoBW add <# downloadBW; '\n' #>
  exit

<# env.setResult("activate-profile", name) #>
<# env.setResult("output-stat-clacl", name) #>

<# endtmp1 #>
```

Sample Owner ID

Step-by-Step Procedure [Table 163 on page 727](#) provides the sample RADIUS attributes that you can use to activate the video-on-demand service

Table 163: Sample RADIUS Attributes

Owner	Owner ID	Value
AAA (RADIUS)	Acct-Session-ID (RADIUS attribute 44)	573498

Sample CLI Command

Step-by-Step Procedure To activate a subscriber session with the configured video-on-demand service:

```
host1(config)#service-management owner-session aaa 573498 service-session
"videoMin(4500000, 192.168.23.58)"
```

- Related Documentation**
- [Service Definitions Overview on page 667](#)
 - [Creating Service Definitions on page 668](#)

- [Managing Your Service Definitions on page 671](#)
- [Understanding Service Manager RADIUS Attributes on page 686](#)
- [Understanding RADIUS Accounting for Service Manager on page 697](#)
- *service-management subscriber-session service-session*

Example: Voice-over-IP Service Definition

The following example shows you how to create a voice-over-IP (VoIP) service.

- [Requirements on page 728](#)
- [Overview on page 728](#)
- [Creating a Voice-over-IP Service on page 729](#)

Requirements

This example uses the following software and hardware components:

- JunosE Release 7.1.0 or higher-numbered releases
- E Series router (ERX7xx models, ERX14xx models, the ERX310 router, the E120 router, or the E320 router)
- ASIC-based line modules that support Fast Ethernet or Gigabit Ethernet

Overview

A VoIP service is a session border controller (SBC) media gateway (MG)-based service that has upstream and downstream components.

The IP address and port for both the subscriber and the opposite end of the phone call were originally negotiated with the SBC. The VoIP service learns the IP addresses and ports for both ends of the call, and then specifies that any traffic to either end is put in the voice traffic class.

Creating a Voice-over-IP Service

Step-by-Step Procedure This example illustrates how to create a VoIP service:



NOTE: Commented text explains the parameterized values in the example of the service definition macro file. Each example is followed by examples of RADIUS information and the CLI command that you can use to activate a subscriber service session.

```
!parameterizes source address and port, destination address and port, and protocol type
<# mgFlow(upDA, upDPort, downDA, downDPort, protType) #>

<# uid := app.servicemanager.getUniqueId #>
<# name := "SM-mgFlow-" $ uid #>
<# oname := "SM-O-mgFlow-" $ uid #>

classifier-list <# name #> <# protType #> any <#upDA #> 0.0.0.0 eq <# upDPort; '\n' #>
policy-list <# name; '\n' #>
  classifier-group <# name #> precedence 2000
  traffic-class voice
  forward

classifier-list <# oname #> <# protType #> any <#downDA #> 0.0.0.0 eq <# downDPort; '\n' #>
policy-list <# oname; '\n' #>
  classifier-group <# oname #> precedence 2000
  traffic-class voice
  forward

profile <# name ; '\n' #>
  ip policy input <# name #> statistics enabled merge
  ip policy output <# oname #> statistics enabled merge

<# env.setResult("activate-profile", name) #>

<# endtmpl #>
```

Sample RADIUS Attributes

Step-by-Step Procedure Table 164 on page 729 provides the sample RADIUS attributes that you can use to activate the VoIP service

Table 164: Sample RADIUS Attributes

RADIUS Attribute	Tag	Value
username	none	client1@isp1.com
activate-service	1	mgFlow(10.10.10.10, 1234, 192.168.45.54, 1234, udp)

Sample CLI Command

Step-by-Step Procedure

To activate a subscriber session with the configured VoIP service:

```
host1(config)#service-management subscriber-session client1@isp1.com interface atm 4/0.1 service-session "mgFlow(10.10.10.10, 1234, 192.168.45.54, 1234, udp)"
```

Related Documentation

- [Service Definitions Overview on page 667](#)
- [Creating Service Definitions on page 668](#)
- [Managing Your Service Definitions on page 671](#)
- [Understanding Service Manager RADIUS Attributes on page 686](#)
- [Understanding RADIUS Accounting for Service Manager on page 697](#)
- `service-management subscriber-session service-session`

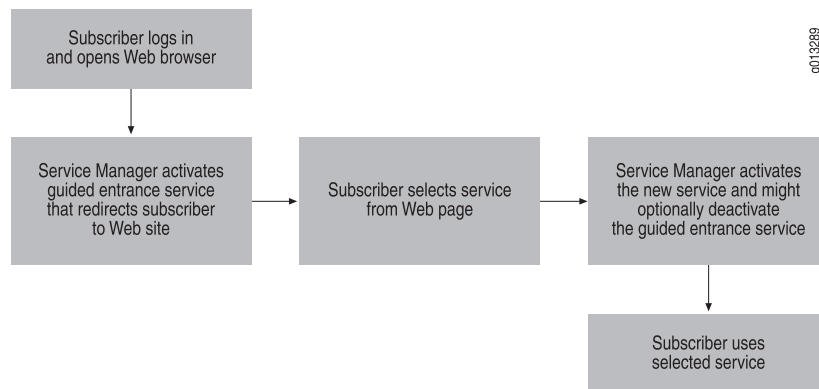
Guided Entrance Service Overview

The guided entrance service enables you to create a controlled Internet browsing environment. Guided entrance-based services, which are sometimes called *walled gardens* or *captive portals*, are becoming increasingly important offerings for service providers. When a subscriber logs in and opens a Web browser, the Service Manager guided entrance service transparently directs the subscriber to a specific website—at that website, the subscriber is presented with a selection of possible services to use. For example, a subscriber might be shown a website that offers services such as:

- **Predefined services**—A group of user-selectable services that meets a variety of needs of a single subscriber. The subscriber might select the high-priced highest access speed to perform critical financial transactions but select a lower speed (and lower cost) service for e-mail. For viewing a real-time sports event, the subscriber can select the video-on-demand service. The subscribers have control over the choice and cost of the services they need and use.
- **Prepaid services**—A group of specific services that have been prepaid by the subscriber. For example, a subscriber who has purchased the sports package service is presented with a webpage that lists the currently available sporting events. Or, a subscriber might prepay a VoIP service for a set amount of time.
- **Controlled-service**—An educational service that enables students at a school to access authorized research sites. Or, a limited service for young children that restricts access to safe, closely monitored, age-appropriate websites.

[Figure 32 on page 731](#) shows the sequence of actions that take place during a guided entrance service.

Figure 32: Guided Entrance



Service Manager requires additional configuration considerations for the guided entrance service.

- The `<# redirectUrlName := "http://" $ serverIp $ ":" $ serverPort #>` command in the service definition—Specifies the HTTP local service to which the subscriber is redirected after login. See [“Example: Guided Entrance Service Definition” on page 731](#) for a sample guided entrance service definition.



NOTE: You must also configure a policy that redirects packets. See *Creating an Exception Rule within a Policy Classifier Group in JunosE Policy Management Configuration Guide* for information on creating redirect policies.

- HTTP local server application—Used by the policy in the activated service to direct a subscriber to a specific website when the subscriber logs in. See [“Configuring the HTTP Local Server to Support Guided Entrance” on page 735](#) for information about the HTTP local server.
- RADIUS Dynamic Request Server and COA messages—Enables RADIUS to dynamically activate the new service that the subscriber selects at the website. Can also optionally deactivate the original guided entrance service session that is used when the subscriber logs in. See the *Configuring RADIUS Dynamic-Request Server* chapter in this guide.

Related Documentation

- [Using COA Messages with Guided Entrance Services on page 734](#)
- [Example: Guided Entrance Service Definition on page 731](#)

Example: Guided Entrance Service Definition

The following example shows you how to create a guided entrance service.

- [Requirements on page 732](#)
- [Overview on page 732](#)
- [Creating a Guided Entrance Service on page 733](#)

Requirements

This example uses the following software and hardware components:

- JunosE Release 7.1.0 or higher-numbered releases
- E Series router (ERX7xx models, ERX14xx models, the ERX310 router, the E120 router, or the E320 router)
- ASIC-based line modules that support Fast Ethernet or Gigabit Ethernet

Overview

When a subscriber logs in and opens a Web browser, the Service Manager guided entrance service transparently directs the subscriber to a specific uniform resource locator (URL) at which the subscriber can choose from a list of available services.

Creating a Guided Entrance Service

Step-by-Step Procedure This example illustrates how to create a guided entrance service:



NOTE: Commented text explains the parameterized values in the example of the service definition macro file. Each example is followed by examples of RADIUS information and the CLI command that you can use to activate a subscriber service session.

```
!parameterizes server address and port
<# http(serverIp, serverPort) #>

<# serviceTag := "http-" #>
<# uid := app.servicemanager.getUniqueId #>
<# genericName := "SM-X-" $ serviceTag $ uid #>
<# genericInputName := "SM-I-" $ serviceTag $ uid #>
<# genericOutputName := "SM-O-" $ serviceTag $ uid #>
<# clacName := genericName #>

<# profileName := genericName #>
<# inputPolicyName := genericInputName #>
<# inputRateLimitName := genericInputName #>
<# outputPolicyName := genericOutputName #>
<# outputRateLimitName := genericOutputName #>

<# exceptionClacName := "exceptionClacPort" $ serverPort #>
<# serverClacName := "serverClacIp" $ serverIp #>
<# redirectUrlName := "http://" $ serverIp $ ":" $ serverPort #>

classifier-list <# serverClacName #> ip any host <# serverIp; '\n' #>

classifier-list <# exceptionClacName #> tcp any any eq <# serverPort; '\n' #>

ip policy-list <# inputPolicyName; '\n' #>
classifier-group <# serverClacName; '\n' #>
    forward
        classifier-group <# exceptionClacName; '\n' #>
            exception http-redirect
        classifier-group *
            filter

profile <# profileName #>
    ip http redirectUrl <# redirectUrlName; '\n' #>
    ip policy input <# inputPolicyName #> statistics enabled merge

<# env.setResult("activate-profile", "" $ profileName) #>

<# endtmp1 #>
```

Sample RADIUS Attributes

Step-by-Step Procedure [Table 165 on page 734](#) provides the sample RADIUS attributes that you can use to activate the guided entrance service

Table 165: Sample RADIUS Attributes

RADIUS Attribute	Tag	Value
username	none	client5@isp1.com
activate-service	1	http(192.168.25.2, 80)

Sample CLI Command

Step-by-Step Procedure To activate a subscriber session with the configured guided entrance service:

```
host1(config)#service-management subscriber-session client5@isp1.com interface atm
5/0.1 service-session "http(192.168.25.2, 80)"
```

Related Documentation

- [Service Definitions Overview on page 667](#)
- [Creating Service Definitions on page 668](#)
- [Managing Your Service Definitions on page 671](#)
- [Understanding Service Manager RADIUS Attributes on page 686](#)
- [Understanding RADIUS Accounting for Service Manager on page 697](#)
- [Guided Entrance Service Overview on page 730](#)
- [Configuring the HTTP Local Server to Support Guided Entrance on page 735](#)
- [Using COA Messages with Guided Entrance Services on page 734](#)
- *service-management subscriber-session service-session*

Using COA Messages with Guided Entrance Services

Typically, a guided entrance service directs a subscriber to a website, where the subscriber can select from a group of available services. When the subscriber selects a new service to use, Service Manager uses a RADIUS COA message to activate the new service—you can also configure RADIUS to deactivate the original guided entrance service. To inform Service Manager to deactivate the original guided entry service, you must include the Deactivate-Service attribute in the RADIUS records of the services that can be selected from the website.

If you configure a guided entrance service, you must also ensure that the router's RADIUS dynamic-request server is enabled and supports COA messages. See the *Configuring RADIUS Dynamic-Request Server* chapter in this guide, for information about the RADIUS dynamic-request server and COA messages.

Table 166 on page 735 describes a partial RADIUS Access-Accept message for a guided entrance service and the COA-Request message for the tiered service that the subscriber subsequently selects from the website. The COA message for the tiered service includes the Deactivate-Service attribute that deactivates the guided entrance service.

Table 166: Deactivating a Guided Entrance Service

Guided Entrance Service Activated at Login		
RADIUS Attribute	Tag	Value
username	none	client5@isp1.com
activate-service	1	http(192.168.25.2, 80)

Tiered Service Selected at Website

RADIUS Attribute	Tag	Value
username	none	client5@isp1.com
activate-service	2	tiered(1280000, 5120000)
deactivate-service		http(192.168.25.2, 80)
service-timeout	2	720
service-statistics	2	2

Related Documentation

- [Guided Entrance Service Overview on page 730](#)
- [Redirection of Subscriber Sessions When HTTP Local Server is Disabled or Not Configured on page 739](#)
- [Configuring the HTTP Local Server to Support Guided Entrance on page 735](#)
- [Example: Guided Entrance Service Definition on page 731](#)

Configuring the HTTP Local Server to Support Guided Entrance

JunosE Software supports an embedded Web server, known as the HTTP local server, which is used to support the Service Manager application's guided entrance service. With guided entrance, subscribers are directed to a specific website when they log in. At the website, the subscribers can then select the service they want to use. You can configure one HTTP local server per virtual router. The HTTP local server is disabled by default.

In lower-numbered releases, the HTTP server listened for and processed only IPv4 exception packets. You can now configure the HTTP local server to listen for and process both IPv4 and IPv6 packets.



NOTE: Currently, the HTTP local server does not support two different ports for IPv4 and IPv6 packets. However, the HTTP local server can listen for both IPv4 and IPv6 exception packets on the same port, simultaneously.

The following topics explain how to configure the HTTP local server to support guided entrance:

- [Configuring the HTTP Local Server to Support Guided Entrance for IPv4 Subscribers on page 736](#)
- [Configuring the HTTP Local Server to Support Guided Entrance for IPv6 Subscribers on page 737](#)

Configuring the HTTP Local Server to Support Guided Entrance for IPv4 Subscribers

To configure the HTTP local server to support guided entrance for IPv4:

1. Access the virtual router context.

```
host1(config)#virtual-router west400
host1:west400(config)#
```

2. Create the HTTP local server.

```
host1:west400(config)#ip http
```

Use the **no** version to delete the HTTP local server.

3. (Optional) Specify a standard IP access list that defines which subscribers can connect to the HTTP local server.

```
host1:west400(config)#ip http access-class chicagoList
```

Use the **no** version to remove the association between the access list and the HTTP local server.

4. (Optional) Specify the port on which the HTTP local server receives connection attempts.

```
host1:west400(config)#ip http port 8080
```

You can specify a port number in the range 1–65535. Use the **no** version to restore the default port number, 80.

5. (Optional) Specify the maximum number of connections that can exist between one IP address and the HTTP local server.

```
host1:west400(config)#ip http same-host-limit 20
```

You can specify a number in the range 0–1000.

Use the **no** version to restore the default number of allowed connections, 3.

6. Specify the maximum time that HTTP local servers maintain connections.


```
host1:west400(config)#ip http max-connection-time 1000
```

You can specify a time in the range 3–7200 seconds, or 0. A value of 0 causes the server to maintain an inactive connection indefinitely. Use the **no** version to restore the default time, 30 seconds.

7. Enable the HTTP local server to listen for and process IPv4 exception packets

```
host1:west400(config)#ip http server
```

Use the **no** version to disable the HTTP local server.

8. Configure the HTTP redirect feature for the profile, interface, or subinterface that will be referenced in the guided entrance service definition.

```
host1:west400(config)#profile guidEnt6
host1:west400(config-profile)#ip http redirectUrl http://ispsite.redirect.com
```

The first access session is typically used by the Service Manager application to provide initial provisioning and service selection for the subscriber. HTTP redirect is per-interface; use the command in Profile Configuration mode for dynamic interfaces; use the command in Interface Configuration mode or Subinterface Configuration mode for static interfaces.

The redirect URL can be a maximum of 230 characters.



NOTE: The HTTP local server must be configured and enabled in the virtual router for the interface on which you use the **ip http redirectUrl** command. Otherwise, the URL redirect operation will fail.

Use the **no** version to restore the default, which disables the HTTP redirect feature.

Configuring the HTTP Local Server to Support Guided Entrance for IPv6 Subscribers

To configure the HTTP local server to support guided entrance for IPv6:

1. Access the virtual router context.

```
host1(config)#virtual-router west400
host1:west400(config)#
```

2. Create the HTTP local server.

```
host1:west400(config)#ipv6 http
```

Use the **no** version to delete the HTTP local server.

3. (Optional) Specify a standard IP access list that defines which subscribers can connect to the HTTP local server.

```
host1:west400(config)#ip http access-class chicagoList
```

Use the **no** version to remove the association between the access list and the HTTP local server.

4. (Optional) Specify the port on which the HTTP local server receives connection attempts.

```
host1:west400(config)#ipv6 http port 8080
```



NOTE: You can modify the port on which the HTTP local server receives connection attempts. However, you must first disable the HTTP local server and then modify the port.

You can specify a port number in the range 1–65535. Use the **no** version to restore the default port number, 80.

5. (Optional) Specify the maximum number of connections that can exist between one IP address and the HTTP local server.

```
host1:west400(config)#ip http same-host-limit 20
```

You can specify a number in the range 0–1000. Use the **no** version to restore the default number of allowed connections, 3.

6. Specify the maximum time that HTTP local servers maintain connections.

```
host1:west400(config)#ip http max-connection-time 1000
```

You can specify a time in the range 3–7200 seconds, or 0. A value of 0 causes the server to maintain an inactive connection indefinitely. Use the **no** version to restore the default time, 30 seconds.

7. Enable the HTTP local server to listen for and process IPv6 exception packets.

```
host1:west400(config)#ipv6 http server
```

Use the **no** version to disable the HTTP local server.

8. Configure the HTTP redirect feature for the IPv6 profile, interface or subinterface to be referenced in the guided entrance service definition.

```
host1:west400(config)#interface gigabitEthernet 6/0
host1:west400(config-if)#ipv6 http redirectUrl http://ispsite.redirect.com
```

The first access session is typically used by the Service Manager application to provide initial provisioning and service selection for the subscriber. HTTP redirect is per-interface; use the command in Interface Configuration mode or Subinterface Configuration mode for static interfaces and use the command in Profile Configuration mode for dynamic interfaces.

The redirect URL can be a maximum of 230 characters.



NOTE: The HTTP local server must be configured and enabled in the virtual router for the interface on which you use the **ipv6 http redirectUrl** command. Otherwise, the URL redirect operation will fail.

Use the **no** version to restore the default, which disables the HTTP redirect feature.

Related Documentation

- [Guided Entrance Service Overview on page 730](#)
- [Using COA Messages with Guided Entrance Services on page 734](#)

- [Example: Guided Entrance Service Definition on page 731](#)
- [Redirection of Subscriber Sessions When HTTP Local Server is Disabled or Not Configured on page 739](#)
- [Configuring the Preservation of the Original URL During Redirection of Subscriber Sessions on page 747](#)
- *ip http*
- *ip http access-class*
- *ip http max-connection-time*
- *ip http port*
- *ip http redirectUrl*
- *ip http same-host-limit*
- *ip http server*
- *ipv6 http*
- *ipv6 http port*
- *ipv6 http redirectUrl*
- *ipv6 http server*

Redirection of Subscriber Sessions When HTTP Local Server is Disabled or Not Configured

The HTTP local server on the router must always be enabled before the subscriber logs in for the redirect URL to be returned. If the HTTP local server is activated on the router after the subscribers have logged in, the URL to which the subscriber's Web browser session needs to be redirected is not returned to the users. Instead, the redirect engine opens a TCP port (8800 by default) and sends an HTTP 302 Found response to the subscriber's browser in response to the request. The subscriber must log out and log in again for the redirection URL to be returned to the subscriber in response to the initial request.

The HTTP redirect URL functionality works correctly only if the HTTP local server is running on the system before subscribers log in. Also, if the HTTP local server is disabled and reenabled, previously logged-in subscribers must log out and reestablish their sessions for the redirect URL to be returned. After you disable and reenable the HTTP local server on the router, the interface configuration details for previously logged-in subscribers are not retained.

Consider a scenario in which two subscribers, subscriber A and subscriber B, are logged in, the HTTP local server for both IP and IPv6 traffic is enabled on the router, and URL redirection functionality is configured. In such a topology, when subscriber A sends an HTTP GET request to the HTTP local server on the router, the subscriber's HTTP session is redirected to the configured URL. If you disable the HTTP local server for IP traffic by using the **no ip http server** command and reenable the HTTP local server by using the

ip http server command, when subscriber B sends an HTTP GET request to the router, the session is not redirected to the configured URL. The user session is redirected to the correct configured site only when subscriber B logs out and logs back in again.

**Related
Documentation**

- [Guided Entrance Service Overview on page 730](#)
- [Configuring the HTTP Local Server to Support Guided Entrance on page 735](#)
- [Example: Guided Entrance Service Definition on page 731](#)

Example: Combined IPv4 and IPv6 Service in a Dual Stack Service Definition

The following example explains how to create a combined IPv4 and IPv6 service in a dual stack.

- [Requirements on page 740](#)
- [Overview on page 740](#)
- [Creating a Combined IPv4 and IPv6 Service on page 742](#)

Requirements

This example uses the following software and hardware components:


- JunosE Release 7.1.0 or higher-numbered releases
- E Series router (ERX7xx models, ERX14xx models, the ERX310 router, the E120 router, or the E320 router)
- ASIC-based line modules that support Fast Ethernet or Gigabit Ethernet

Overview

When you configure a combined IPv4 and IPv6 service in a dual stack, the policies defined in the interface profile are attached to the appropriate interfaces based on the type of the interface. For example, all IPv4 policies are attached to the IPv4 interface and all IPv6 policies are attached to the IPv6 interface.


[Figure 33 on page 741](#) shows a topology in which the C-VLAN interface on the customer edge device is connected to the ingress IPv4 and IPv6 interfaces on the provider edge or E Series router. A combined IPv4/IPv6 service, which contains a hierarchical policy and an external parent group with a rate-limit profile that is associated with the hierarchical policy, is applied at the secondary input stage on the router. The incoming voice-over-IP classified traffic flows for IPv4 and IPv6 subscribers are allocated a total of 64 Kbps. The common rate limit cannot drop voice-over-IP packets, but must limit the total flow (for IPv4 and IPv6 interfaces) to 64 Kbps.

Figure 33: Input Traffic Flow with Rate-Limit Profile on an External Parent Group for a Combined IPv4/IPv6 Service



Similarly, for traffic flowing from the provider edge device to the C-VLAN interface for voice-over-IP packets, [Figure 34 on page 741](#) shows how the rate-limit profile in the external parent group associated with a hierarchical policy parameter applied to the egress IPv4 and IPv6 interfaces limits the voice-over-IP traffic flowing to the C-VLAN interface on the customer edge device.

Figure 34: Output Traffic Flow with Rate-Limit Profile on an External Parent Group for a Combined IPv4/IPv6 Service



Creating a Combined IPv4 and IPv6 Service

Step-by-Step Procedure

The following example shows the service macro definition file that creates a voice-over-IP service for the topology described above.

```
<# combined_service(inBw, outBw, VBG1, VB6G1, NODE) #>

<# uid := app.servicemanager.getUniqueId #>
<# genericName := "combined-service-" $ uid #>
<# SACLacIName := genericName $ "SA" #>
<# profileName := genericName #>

policy-parameter v4v6-<# uid #> hierarchical
  aggregation-node <# NODE #><# '\n' #>

rate-limit-profile rlpv4v6-<# genericName #>-vb-out one-rate hier
committed-rate <# outBw #><# '\n' #>
committed-action transmit unconditional
conformed-action transmit unconditional

rate-limit-profile rlpv4v6-<# genericName #>-vb-in one-rate hier
committed-rate <# inBw #><# '\n' #>
committed-action transmit unconditional
conformed-action transmit unconditional

parent-group vb-v4v6-<# uid #>-in
  rate-limit-profile rlpv4v6-<# genericName #>-vb-in

parent-group vb-v4v6-<# uid #>-out
  rate-limit-profile rlpv4v6-<# genericName #>-vb-out

classifier-list cl46-4-<# genericName #>-vb-in ip any host <# VBG1 #> <# '\n' #>
classifier-list cl46-4-<# genericName #>-vb-out ip host <# VBG1 #> any

ipv6 classifier-list cl46-6-<# genericName #>-vb-in destination-host <# VB6G1 #><# '\n' #>
ipv6 classifier-list cl46-6-<# genericName #>-vb-out source-host <# VB6G1 #><# '\n' #>

ip policy-list pl-v4v6-<# genericName #>-in
  classifier-group cl46-4-<# genericName #>-vb-in external parent-group vb-v4v6-<# uid #>-in parameter
v4v6-<# uid #><# '\n' #>
  forward

ip policy-list pl-v4v6-<# genericName #>-out
  classifier-group cl46-4-<# genericName #>-vb-out external parent-group vb-v4v6-<# uid #>-out parameter
v4v6-<# uid #><# '\n' #>
  forward

ipv6 policy-list pl6-v4v6-<# genericName #>-in
  classifier-group cl46-6-<# genericName #>-vb-in external parent-group vb-v4v6-<# uid #>-in parameter
v4v6-<# uid #><# '\n' #>
  forward

ipv6 policy-list pl6-v4v6-<# genericName #>-out
  classifier-group cl46-6-<# genericName #>-vb-out external parent-group vb-v4v6-<# uid #>-out parameter
v4v6-<# uid #><# '\n' #>
  forward
```

```
profile <# profileName #><# '\n' #>
  ip policy output p1-v4v6-<# genericName #>-out statistics enabled merge
  ip policy secondary-input p1-v4v6-<# genericName #>-in statistics enabled merge
  ipv6 policy output p16-v4v6-<# genericName #>-out statistics enabled merge
  ipv6 policy secondary-input p16-v4v6-<# genericName #>-in statistics enabled merge

<# env.setResult("activate-profile", profileName) #>
<# env.setResult("service-interface-type", "ipv4-ipv6") #>
<# env.setResult("secondary-input-stat-epg", "vb-v4v6-"$ uid $"-in v4v6-"$ uid "$") #>
<# env.setResult("output-stat-epg", "vb-v4v6-"$ uid $"-out v4v6-"$ uid "$") #>
<# endtmpl #>
```

In the service definition macro, a hierarchical policy parameter for the rate limit is created with an aggregation node value. The aggregation node stores a single rate-limit instance and statistics for this rate-limit. An external rate-limit aggregation node can be defined by the 4-tuple (slot, direction, external parent group name, parameter value). Each reference to a policy parameter in a policy is substituted with its value for all attachments of this policy at the interface.

Two rate-limit profiles are created, one each for the ingress and egress interfaces. Rate limiters are implemented using a dual token bucket scheme: a token bucket for conformed (yellow) packets and a token bucket for committed (green) packets. The following are the attributes configured in the rate-limit profile applied to ingress and egress interfaces:

- The committed rate for the rate-limit profile is entered as a specified value.
- The committed action, which specifies the action for packets conforming to the committed rate and committed burst size and conforming to the exceed rate and exceed burst size for a rate-limit profile is set to receive transmit unconditional.
- The conformed action, which sets the action for packets not conforming to the committed rate and committed burst size, but conforming to the peak rate and peak burst size for a rate-limit profile is set to receive transmit unconditional.

Two external parent groups, one each for the ingress and egress interfaces, that reference the rate-limit profiles created for incoming and outgoing traffic, are created and specified in the service definition.

Classifier control lists for ingress IPv4 and IPv6 traffic, and for egress IPv4 and IPv6 traffic, are also created. These classifiers classify traffic based on source and destination addresses.

The input and output classifier lists for IPv4 traffic are used in IP policy lists that are attached to the ingress and egress IPv4 interfaces respectively. The input and output classifier lists for IPv6 traffic are used in IPv6 policy lists that are attached to the ingress and egress IPv6 interfaces respectively. The **external parent-group** keyword creates an external parent group in a rate-limit hierarchy for IPv4 and IPv6. All packets matching the classifier are sent to the parent group for further processing.

The policy lists for voice-over-IP traffic are configured in the service definition macro file that creates a combined IPv4/IPv6 service to be applied to the ingress IPv4 and IPv6 interfaces.

A profile is created that you want to attach to the service session. The IPv4 and IPv6 policies for voice-over-IP traffic arriving at the IPv4 and IPv6 interfaces respectively are applied to the secondary input stage. The IPv4 and IPv6 policies for voice-over-IP traffic leaving the IPv4 and IPv6 interfaces respectively are applied to the output stage. Statistics collection is enabled for the policies referenced in the service macro using the **statistics enabled** keyword in the command used for policy attachment in the profile. The **merge** keyword enables merging of multiple policies to form a single policy.

The `<# env.setResult("activate-profile", profileName) #>` command specifies the interface profile to be used on activation of the interface. After the subscriber service session is activated, the policies defined in the interface profile are applied to both IPv4 and IPv6 interfaces. The elements in the profile to be attached to the interfaces are determined by the type of the interface.

The `<# env.setResult("service-interface-type", "ipv4-ipv6") #>` command configures the service macro to be used for IPv4 and IPv6 interfaces in a dual stack. The profile identifier returned from the activate-profile object will be applied to both IPv4 and IPv6 interfaces.

The service definition macro is configured to collect input and output statistics associated with external parent groups in a hierarchical policy for IPv4 and IPv6 subscribers as follows:

```
<# env.setResult("secondary-input-stat-epg", "vb-v4v6-"$ uid $"-in v4v6-"$ uid $"") #>
<# env.setResult("output-stat-epg", "vb-v4v6-"$ uid $"-out v4v6-"$ uid $"") #>
```

The `<# env.setResult("secondary-input-stat-epg", "vb-v4v6-"$ uid $"-in v4v6-"$ uid $"") #>` command specifies that Service Manager track statistics associated with the external parent group named vb-v4v6-in and the corresponding hierarchical policy named v4v6, and that this external parent group is associated with the policy that is attached at the input stage.

The `<# env.setResult("output-stat-epg", "vb-v4v6-"$ uid $"-out v4v6-"$ uid $"") #>` command specifies that Service Manager track the output statistics associated with the external parent group named vb-v4v6-out and the corresponding hierarchical policy named v4v6, which is associated with the policy attached at the output stage.

The input and output statistics associated with the external parent group are collected and forwarded to the Service Manager to be displayed in the Acct-Stop and Interim-Acct messages.

If you use the **secondary-input-stat-clacl** and **output-stat-clacl** objects in the service macro to track Service Manager statistics, the values returned in the output of the **show service-management** command do not accurately reflect the packets that are rate-limited. In this case, although some of the packets that were classified by the classifier lists are dropped by the rate-limiter on the external parent group, the Service Manager statistics collection application counts all the packets that were classified without excluding those that were dropped by the rate limiter. As a result, the values returned by the **output-stat-clacl** and **secondary-input-stat-clacl** objects represent more packets than those sent to the subscriber and core interfaces respectively.

Using the macro that has been described here, you can configure the following combined service, for example:

combined_service(64000, 64000, 10.0.0.1, 2001::1, vlan)

where

- 64000—Bandwidth for outbound traffic, denoted as *outBw* in the macro
- 64000—Bandwidth for inbound traffic, denoted as *inBw* in the macro
- 10.0.0.1—Host IP address for IPv4 subscribers, denoted as *VBG1* in the macro
- 2001::1—Host IP address for IPv6 subscribers, denoted as *VB6G1* in the macro
- vlan—Interface on which the service is configured, denoted as *NODE* in the macro

**Related
Documentation**

- [Combined and Independent IPv4 and IPv6 Services in a Dual Stack Overview on page 694](#)
- [Activation and Deactivation of IPv4 and IPv6 Services in a Dual Stack on page 696](#)
- [Service Interim Accounting for IPv4 and IPv6 Services in a Dual Stack Overview on page 705](#)
- [Configuring the HTTP Local Server to Support Guided Entrance on page 735](#)

Preservation of the Original URL During Redirection of Subscriber Sessions

When guided entrance service is used and a digital subscriber line (DSL) user logs in for the first time and opens a Web browser, the subscriber is directed to a specific URL for user provisioning. Provisioning involves application of a service definition which enables all HTTP traffic to be handled temporarily by the HTTP local server. After the provisioning process, the service definition is removed and traffic no longer flows to the HTTP local server. At this point, the user is redirected to the original requested URL. The HTTP local server redirect feature supports the preservation of the original URL as a variable in the redirect URL. If preservation of the original URL feature is enabled, then the user's HTTP session is directed back to the URL which was requested before the redirection.

**Related
Documentation**

- [Configuring the Preservation of the Original URL During Redirection of Subscriber Sessions on page 747](#)

Configuring the Preservation of the Original URL During Redirection of Subscriber Sessions

You can configure the HTTP local server to preserve the URL, originally requested by a user, as a variable in the redirect URL.

- From Profile Configuration mode, enable preservation of the original URL on the HTTP local server.

- To configure preservation of the original URL on the IPv4 profile:

```
host1(config-profile)#ip http redirectUrl
"http://ispite.redirect.com/accessDenied.do?url=%(url)" preserveOriginalUrl
```

- To configure preservation of the original URL on the IPv6 profile:

```
host1(config-profile)#ipv6 http redirectUrl
"http://ispite.redirect.com/accessDenied.do?url=%(url)" preserveOriginalUrl
```



NOTE: You must press Ctrl+v before typing “?” in the CLI. You must also ensure that the redirect URL is of the HTTP query type according to the server language supported by the redirect server.

Related Documentation

- [Preservation of the Original URL During Redirection of Subscriber Sessions on page 746](#)
- *ip http redirectUrl*
- *ipv6 http redirectUrl*

CHAPTER 29

Monitoring Service Manager

This chapter describes how to monitor the Service Manager application. This chapter discusses the following topics:

- [Setting a Baseline for HTTP Local Server Statistics on page 749](#)
- [Monitoring the Connections to the HTTP Local Server on page 750](#)
- [Monitoring the Configuration of the HTTP Local Server on page 750](#)
- [Monitoring Statistics for Connections to the HTTP Local Server on page 751](#)
- [Monitoring Profiles for the HTTP Local Server on page 752](#)
- [Monitoring the Default Interval for Interim Accounting of Services on page 753](#)
- [Monitoring the Status of the Service Manager License on page 753](#)
- [Monitoring Profiles for Service Manager on page 754](#)
- [Monitoring IPv4 and IPv6 Interfaces for Service Manager on page 755](#)
- [Monitoring Service Definitions on page 765](#)
- [Monitoring Service Session Profiles on page 766](#)
- [Monitoring Active Owner Sessions with Service Manager on page 767](#)
- [Monitoring Active Subscriber Sessions with Service Manager on page 770](#)
- [Monitoring the Number of Active Subscriber and Service Sessions with Service Manager on page 773](#)
- [Verifying Computation of Service Session Accounting Based on Scheduler Profiles on page 773](#)

Setting a Baseline for HTTP Local Server Statistics

You can set a baseline for HTTP server statistics.

The system implements the baseline by reading and storing the statistics at the time the baseline is set and then subtracting this baseline whenever baseline-relative statistics are retrieved.

To set a baseline:

- Include the **baseline ip http** command at the User Exec or Privilege Exec level:
`host1#baseline ip http`

There is no **no** version.

- Related Documentation**
- [Monitoring Statistics for Connections to the HTTP Local Server on page 751](#)
 - *baseline ip http*

Monitoring the Connections to the HTTP Local Server

Purpose Display information about the connections to the HTTP local server.

Action To display information about the HTTP local server:

```
host1#show ip http scalar
Maximum connection length: 1000 seconds
Current number of http servers: 5
Number of enabled http servers: 2
Current number of http connections: 15
Peak number of http connections: 125
Maximum number of http connections: 1000
```

Meaning [Table 167 on page 750](#) lists the **show ip http scalar** command output fields.

Table 167: show ip http scalar Output Fields

Field Name	Field Description
Maximum connection length	Maximum time that the HTTP local server maintains an inactive connection, in seconds
Current number of http servers	Number of configured Web servers
Number of enabled http servers	Number of Web servers enabled
Current number of http connections	Number of connections from subscribers to HTTP local servers
Peak number of http connections	Highest number of connections from subscribers to HTTP local servers
Maximum number of http connections	Maximum number of connections allowed from subscribers to HTTP local servers

- Related Documentation**
- *show ip http scalar*

Monitoring the Configuration of the HTTP Local Server

Purpose Display information about the configuration of the HTTP local server.

Action To display information about the HTTP local server:

```
host1#show ip http server
Admin status: enabled
Access class: not defined
```

Listening port: 80
 Same host limit: 3
 Protocol: IPv6

Meaning [Table 168 on page 751](#) lists the **show ip http server** command output fields.

Table 168: show ip http server Output Fields

Field Name	Field Description
Admin status	Status of the HTTP local server in the software: enabled or disabled. Enabled implies that the HTTP local server can listen for IPv4, IPv6, or both IPv4 and IPv6 exception packets.
Access class	Name of a standard IP access list that determines which hosts can log on to the HTTP local server
Listening port	Port that the HTTP local server uses to track requests for connection
Same host limit	Maximum number of connections allowed between one IP address and the HTTP local server
Protocol	Protocols that the HTTP local server is listening for: IPv4, IPv6, or IPv4 and IPv6.

Related Documentation

- *show ip http server*

Monitoring Statistics for Connections to the HTTP Local Server

Purpose Display statistics about the connections to the HTTP local server.

Action To display statistics about HTTP local server with the baseline values subtracted:

```
host1#show ip http statistics delta
Server enable count: 1
Server disable count: 0
Same host enforced: 0
Access class denies: 0
No resource failures: 0
Http connections created: 2
Http connections terminated: 2
Http connections aged out: 1
Urls successfully served: 0
Malformed http requests: 0
Urls not found: 0
```

Meaning [Table 169 on page 752](#) lists the **show ip http statistics** command output fields.

Table 169: show ip http statistics Output Fields

Field Name	Field Description
Server enable count	Total number of enabled HTTP local servers
Server disable count	Total number of disabled HTTP local servers
Same host enforced	Number of connections dropped because the limit for connections from one IP address to the HTTP local server was exceeded
Access class denies	Number of connections dropped because of a problem with the standard IP access list that defines the hosts that can access the HTTP local server
No resource failures	Number of connections dropped because of system memory limitations
Http connections created	Total number of HTTP connections established
Http connections terminated	Total number of HTTP connections ended
Http connections aged out	Total number of HTTP connections that expired because they exceeded the maximum allowed connection time
Urls successfully served	Total number of webpages displayed
Malformed http requests	Number of HTTP requests that failed because the format was incorrect
Urls not found	Number of webpages not found

Related Documentation

- *show ip http statistics*

Monitoring Profiles for the HTTP Local Server

Purpose Display information about the redirect URL used for guided entrance services.

Action To display information about the redirect URL used by the HTTP local server:

```
host1#show profile name guidedProfile2
Profile                               : guidedProfile2
.
.
.
Auto Detect                          : Disabled
Auto Configure                       : Disabled
IP FlowStats                         : Disabled
Ip http redirect Url : myredirect.html
```


Ipv6 http redirect Url: myredirect.html

Meaning [Table 170 on page 753](#) lists the **show profile** command output fields.

Table 170: show profile Output Fields

Field Name	Field Description
Ip http redirect Url	URL of the webpage used for Service Manager guided entrance services for IPv4
Ipv6 http redirect Url	URL of the webpage used for Service Manager guided entrance services for IPv6

Related Documentation

- *show profile*

Monitoring the Default Interval for Interim Accounting of Services

Purpose Display the default interval used for interim accounting for services associated with users on the virtual router. An entry of 0 indicates that the feature is disabled.

Action To display the default interval used for interim accounting:

```
host1:vrXyz7#show aaa service accounting interval
service-acct-interval 60
```

Meaning [Table 171 on page 753](#) lists the **show aaa service accounting interval** command output fields.

Table 171: show aaa service accounting interval Output Fields

Field Name	Field Description
service-acct-interval	Value of the default interval

Related Documentation

- *show aaa service accounting interval*

Monitoring the Status of the Service Manager License

Purpose Display the status of the Service Manager license.

Action To display the status of the Service Manager license:

```
host1#show license service-management
service management license is set
```

Meaning [Table 172 on page 754](#) lists the **show license service-management** command output fields.

Table 172: show license service-management Output Fields

Field Name	Field Description
service management license	Status of the license: set (enabled) or not set (disabled)

Related Documentation

- *show license service-management*

Monitoring Profiles for Service Manager

Purpose Display information about the policies and QoS configurations referenced in profiles.

Action To display information about a specific profile:

host1#show profile name video

IP Output Policy : video statistics disabled

IP Secondary Input Policy : video statistics disabled

qos-parameter vidburst 1000

qos-parameter vidrate 500000

qos-profile vid512k

host1#show profile name foo

IP Policy Parameter foo : 100000 increase, reference rate

IP Input Policy : p1 statistics disabled

ERX-00-16-c2#show profile name p2

IP Policy Parameter foo : 100000, reference rate

IP Input Policy : p1 statistics disabled

To display a list of profiles configured on the router:

host1#show profile brief

Meaning [Table 173 on page 754](#) lists the **show profile** command output fields.

Table 173: show profile Output Fields

Field Name	Field Description
Input Policy	Name of input policy and whether statistics are enabled or disabled
Output Policy	Name of output policy and whether statistics are enabled or disabled
qos-parameter	Name and value of the QoS parameter assigned to the profile
qos-profile	Name of the QoS profile assigned to the profile

Related • *show profile*
Documentation

Monitoring IPv4 and IPv6 Interfaces for Service Manager

Purpose Display status information about the IP and IPv6 interfaces.

Action To display information about a specific IP interface.

```
host1#show ip interface gigabitEthernet 1/1.200
GigabitEthernet1/1 line protocol Ethernet is up, ip is not present
  Network Protocols: IP
  Multipath mode = hashed
  Auto Configure = disabled
  Auto Detect = disabled
  Inactivity Timer = disabled
  Use Framed Routes = disabled
  ARP spoof checking = disabled
  Warm-restart initial-sequence-preference: Operational = 0 Administrative = 0

  In Received Packets 0, Bytes 0
    Unicast Packets 0, Bytes 0
    Multicast Packets 0, Bytes 0
  In Policed Packets 0, Bytes 0
  In Error Packets 0
  In Invalid Source Address Packets 0
  In Discarded Packets 0
  Out Forwarded Packets 0, Bytes 0
    Unicast Packets 0, Bytes 0
    Multicast Routed Packets 0, Bytes 0
  Out Scheduler Dropped Packets 0, Bytes 0
  Out Policed Packets 0, Bytes 0
  Out Discarded Packets 0

  queue 0: traffic class best-effort, bound to ip GigabitEthernet1/1
    Queue length 0 bytes
    Forwarded packets 0, bytes 0
    Dropped committed packets 0, bytes 0
    Dropped conformed packets 0, bytes 0
    Dropped exceeded packets 0, bytes 0

Http Redirect Url: http://www.juniper.net
```

To display information about a specific IPv6 interface.

```
host1#show ipv6 interface FastEthernet 9/0.6
FastEthernet9/0.6 line protocol VlanSub is up, ipv6 is up
  Description: IPv6 interface in Virtual Router Hop6
  Network Protocols: IPv6
  Link local address: fe80::90:1a00:740:31cd
  Internet address: 6:1::1/64
  Operational MTU 1500 Administrative MTU 0
  Operational speed 1000000000 Administrative speed 0
  Creation type Static
  ND reachable time is 3600000 milliseconds
  ND duplicate address detection attempts is 100
  ND neighbor solicitation retransmission interval is 1000 milliseconds
  ND proxy is enabled
  ND RA source link layer is advertised
  ND RA interval is 200 seconds, lifetime is 1800 seconds
```

```
ND RA managed flag is disabled, other config flag is disabled
ND RA advertising prefixes configured on interface

In Received Packets 0, Bytes 0
  Unicast Packets 0, Bytes 0
  Multicast Packets 0, Bytes 0
In Total Dropped Packets 0, Bytes 0
  In Policed Packets 0
  In Invalid Source Address Packets 0
  In Error Packets 0
  In Discarded Packets 0

Out Forwarded Packets 8, Bytes 768
  Unicast Packets 8, Bytes 768
  Multicast Routed Packets 0, Bytes 0
Out Total Dropped Packets 5, Bytes 0
  Out Scheduler Dropped Packets 0, Bytes 0
  Out Policed Packets 0
  Out Discarded Packets 5

queue 0: traffic class best-effort, bound to ipv6 FastEthernet9/0.6
  Queue length 0 bytes
  Forwarded packets 0, bytes 0
  Dropped committed packets 0, bytes 0
  Dropped conformed packets 0, bytes 0
  Dropped exceeded packets 0, bytes 0

IPv6 policy input ipv6InPol25
  rate-limit-profile Rlp2Mb classifier-group clgA entry 1
    Committed: 0 packets, 0 bytes
    Conformed: 0 packets, 0 bytes
    Exceeded: 0 packets, 0 bytes
  rate-limit-profile Rlp8Mb
    Committed: 0 packets, 0 bytes
    Conformed: 0 packets, 0 bytes
    Exceeded: 0 packets, 0 bytes
IPv6 policy output ipv6PolOut2
  rate-limit-profile RlpOutA classifier-group clgB entry 1
    Committed: 0 packets, 0 bytes
    Conformed: 0 packets, 0 bytes
    Exceeded: 0 packets, 0 bytes
  rate-limit-profile RlpOutB
    Committed: 0 packets, 0 bytes
    Conformed: 0 packets, 0 bytes
    Exceeded: 0 packets, 0 bytes
IPv6 policy local-input ipv6PolLocIn5
  rate-limit-profile Rlp1Mb classifier-group clgC entry 1
    Committed: 0 packets, 0 bytes
    Conformed: 0 packets, 0 bytes
    Exceeded: 0 packets, 0 bytes
  rate-limit-profile Rlp5Mb
    Committed: 0 packets, 0 bytes
    Conformed: 0 packets, 0 bytes
    Exceeded: 0 packets, 0 bytes
queue 0: traffic class best-effort, bound to ipv6 FastEthernet9/0.6
  Queue length 0 bytes
  Forwarded packets 0, bytes 0
  Dropped committed packets 0, bytes 0
  Dropped conformed packets 0, bytes 0
  Dropped exceeded packets 0, bytes 0
Http Redirect Url: http://www.juniper.net
```

Meaning Table 174 on page 757 lists the **show ip interface** command output fields.

Table 174: show ip interface Output Fields

Field Name	Field Description
interface	Interface type and specifier.
interface status	Status of the interface.
HTTP Redirect Url	URL to which a subscriber's initial web browser session is redirected
line protocol	Status of the line protocol.
Description	Text description or alias if configured for the interface
Link up/down trap	Status of SNMP link up/down traps on the interface
Internet Address	IP address of the interface
IP Statistics Rcvd	
local destination	Frames with this router as destination
hdr errors	Number of packets containing header errors
addr errors	Number of packets containing addressing errors
unkn proto	Number of packets received containing unknown protocols
discards	Number of discarded packets
IP Statistics Frags	
reasm ok	Number of reassembled packets
reasm req	Number of requests for reassembly
reasm fails	Number of reassembly failures
frag ok	Number of packets fragmented successfully
frag req	Number of frames requiring fragmentation
frag fails	Number of packets unsuccessfully fragmented
IP Statistics Sent	
generated	Number of packets generated

Table 174: show ip interface Output Fields (*continued*)

Field Name	Field Description
no routes	Number of packets that could not be routed
discards	Number of packets that could not be routed and were discarded
ICMP Statistics Rcvd	
errors	Error packets received
dst unreach	Packets received with destination unreachable
time excd	Packets sent with time-to-live exceeded
param probs	Packets sent with parameter errors
src quench	Source quench packets sent
redirect	Send packets redirect
timestamp req	Requests for a timestamp
timestamp rpy	Replies to timestamp requests
addr mask req	Address mask requests
addr mask rpy	Address mask replies
ICMP Statistics Sent	
errors	Error packets received
dst unreach	Packets received with destination unreachable
time excd	Packets sent with time-to-live exceeded
param probs	Packets sent with parameter errors
src quench	Source quench packets sent
redirect	Send packets redirect
timestamp req	Requests for a timestamp
timestamp rpy	Replies to timestamp requests
addr mask req	Address mask requests

Table 174: show ip interface Output Fields (*continued*)

Field Name	Field Description
addr mask rpy	Address mask replies
ARP spoof checking	<p>Status of the check for spoofed ARP packets received on an IP interface. Possible states: enabled or disabled.</p> <p>NOTE: This field is not displayed when you use the detail keyword.</p>
In Received Packets, Bytes	Total number of packets and bytes received on the IP interface.
Unicast Packets, Bytes	Unicast packets and bytes received on the IP interface; link-local received multi-cast packets (non-multicast-routed frames) are counted as unicast packets.
Multicast Packets, Bytes	Multicast packets and bytes received on the IP interface which are then multicast-routed are counted as multicast packets.
In Forwarded Packets, Bytes	Packets and bytes forwarded into an output IP interface
In Total Dropped Packets, Bytes	Total number of packets and bytes that were dropped on the interface
In Policed Packets	Packets discarded on a receive IP interface for any of the following reasons: exceeding the token bucket limit, exceeding the rate limit, a drop action in a policy, discarded MAC validation packets, a destination address lookup failure or when the destination address is an IP interface that has a route configured to the null interface.
In Invalid Source Address Packets	Packets discarded on a receive IP interface because of invalid IP source address
In Error Packets	Packets discarded on a receive IP interface because of IP header errors
In Discarded Packets	Packets discarded on the ingress interface because of a configuration problem rather than a problem with the packet itself
In Fabric Dropped Packets	Packets discarded on a receive IP interface because of internal fabric congestion
Out Forwarded Packets, Bytes	Total number of packets and bytes forwarded out of the IP interface

Table 174: show ip interface Output Fields (*continued*)

Field Name	Field Description
Unicast Packets, Bytes	Unicast packets and bytes forwarded out of the IP interface
Multicast Routed Packets, Bytes	Multicast packets and bytes forwarded out of the IP interface
Out Requested Packets, Bytes	Packets and bytes requested to be forwarded out an IP interface
Out Total Dropped Packets, Bytes	Total number of packets and bytes that were discarded on the egress interface
Out Scheduler Drops Committed Packets, Bytes	Packets and bytes dropped by the scheduler even though they had a committed traffic contract
Out Scheduler Drops Conformed Packets, Bytes	Packets and bytes dropped by the scheduler even though they conformed to the traffic contract
Out Scheduler Drops Exceeded Packets, Bytes	Packets and bytes dropped by the scheduler because they exceeded the contract
Out Policed Packets	Packets discarded on the egress interface because of rate limiting
Out Discarded Packets	Packets discarded on the egress interface because of a configuration problem rather than a problem with the packet itself
Out Fabric Dropped Packets	Packets dropped because of internal fabric congestion

Table 175 on page 760 lists the **show ipv6 interface** command output fields.

Table 175: show ipv6 interface Output Fields

Field Name	Field Description
Description	Text description or alias if configured for the interface
HTTP Redirect Url	Url to which a subscriber's initial web browser session is redirected
Internet Address	IP address of the interface
Link local address	Local IPv6 address of this interface
Network Protocols	Network protocols configured on this interface
IPv6 Statistics Rcvd	

Table 175: show ipv6 interface Output Fields (*continued*)

Field Name	Field Description
local destination	Frames with this router as destination
hdr errors	Number of packets containing header errors
addr errors	Number of packets containing addressing errors
unkn proto	Number of packets received containing unknown protocols
discards	Number of discarded packets
IP Statistics Sent	
generated	Number of packets generated
no routes	Number of packets that could not be routed
discards	Number of packets that could not be routed and were discarded
ICMPv6 Statistics Rcvd	
total	Total number of received packets
errors	Error packets received
destination unreachable	Packets received with destination unreachable
admin unreachable	Packets received because the destination was administratively unreachable (for example, the packet encountered a firewall filter)
param probs	Packets sent with parameter errors
time excd	Packets sent with time-to-live exceeded
pkt too big	Number of packet-too-big messages received that indicate a packet was too large to forward because of the allowed MTU size
redirects	Received packet redirects
echo requests	Echo request (ping) packets
echo replies	Echo replies received
rtr solicits	Number of received router solicitations

Table 175: show ipv6 interface Output Fields (*continued*)

Field Name	Field Description
rtr advertisements	Number of received router advertisements
neighbor solicits	Number of received neighbor solicitations
neighbor advertisements	Number of received neighbor advertisements
Group membership (queries, responses, reductions)	Number of queries, responses, and reduction requests received from within a group to which the interface is assigned
ICMPv6 Statistics Sent	
total	Total number of received packets
errors	Error packets received
destination unreachable	Packets received with destination unreachable
admin unreachable	Packets received because the destination was administratively unreachable (for example, the packet encountered a firewall filter)
param probs	Packets sent with parameter errors
time excd	Packets sent with time-to-live exceeded
pkt too big	Number of packet-too-big messages received that indicate a packet was too large to forward because of the allowed MTU size
redirects	Received packet redirects
echo requests	Echo request (ping) packets
echo replies	Echo replies received
rtr solicits	Number of received router solicitations
rtr advertisements	Number of received router advertisements
neighbor solicits	Number of received neighbor solicitations
neighbor advertisements	Number of received neighbor advertisements
Group membership (queries, responses, reductions)	Number of queries, responses, and reduction requests received from within a group to which the interface is assigned

Table 175: show ipv6 interface Output Fields (*continued*)

Field Name	Field Description
Operational MTU	Value of the MTU
Administrative MTU	Value of the MTU if it has been administratively overridden using the configuration
Operational speed	Speed of the interface
Administrative speed	Value of the speed if it has been administratively overridden using the configuration
Creation type	Method by which the interface was created (static or dynamic)
ND reachable time	Amount of time (in milliseconds) that the neighbor is expected to remain reachable
ND duplicate address detection attempts	Number of times that the router attempts to determine a duplicate address
ND neighbor solicitation retransmission interval	Interval in which the router retransmits neighbor solicitations
ND proxy	Indicates whether the router will reply to solicitations on behalf of a known neighbor
ND RA source link layer	Indicates whether the RA includes the link layer
ND RA interval	Interval (in seconds) of the neighbor discovery router advertisement
ND RA lifetime	Lifetime (in seconds) of the neighbor discovery router advertisement
ND RA managed flag	State of the neighbor discovery router advertisement managed flag
ND RA other config flag	State of the neighbor discovery router advertisement other config flag
ND RA advertising prefixes	Configured advertisement prefixes for neighbor discovery router advertisement
In Received Packets, Bytes	Total number of packets and bytes received on the IP interface.
Unicast Packets, Bytes	Unicast packets and bytes received on the IP interface; link-local received multi-cast packets (non-multicast-routed frames) are counted as unicast packets.

Table 175: show ipv6 interface Output Fields (*continued*)

Field Name	Field Description
Multicast Packets, Bytes	Multicast packets and bytes received on the IP interface which are then multicast-routed are counted as multicast packets.
In Total Dropped Packets, Bytes	Total number of packets and bytes that were dropped on the interface
In Policed Packets	Packets discarded on a receive IP interface for any of the following reasons: exceeding the token bucket limit, exceeding the rate limit, a drop action in a policy, discarded MAC validation packets, a destination address lookup failure or when the destination address is an IP interface that has a route configured to the null interface.
In Invalid Source Address Packets	Packets discarded on a receive IP interface because of invalid IP source address
In Error Packets	Packets discarded on a receive IP interface because of IP header errors
In Discarded Packets	Packets discarded on the ingress interface because of a configuration problem rather than a problem with the packet itself
Out Forwarded Packets, Bytes	Total number of packets and bytes forwarded out of the IP interface
Unicast Packets, Bytes	Unicast packets and bytes forwarded out of the IP interface
Multicast Routed Packets, Bytes	Multicast packets and bytes forwarded out of the IP interface
Out Total Dropped Packets, Bytes	Total number of packets and bytes that were discarded on the egress interface
Out Scheduler Dropped Packets, Bytes	Number of outbound packets and bytes dropped by the scheduler
Out Policed Packets	Packets discarded on the egress interface because of rate limiting
Out Discarded Packets	Packets discarded on the egress interface because of a configuration problem rather than a problem with the packet itself
IPv6 policy	Type (input, output, local-input) and name of policy
rate-limit-profile	Name of profile

Table 175: show ipv6 interface Output Fields (*continued*)

Field Name	Field Description
classifier-group entry	Entry index
Committed	Number of packets and bytes conforming to the committed access rate
Conformed	Number of packets and bytes that exceed the committed access rate but conform to the peak access rate
Exceeded	Number of packets and bytes exceeding the peak access rate
queue, traffic class, bound to ipv6	Queue and traffic class bound to the specified IPv6 interface
Queue length	Number of bytes in queue
Dropped committed packets, bytes	Total number of committed packets and bytes dropped by this interface
Dropped conformed packets, bytes	Total number of conformed packets and bytes dropped by this interface
Dropped exceeded packets, bytes	Total number of exceeded packets and bytes dropped by this interface

- Related Documentation**
- *show ip interface*
 - *show ipv6 interface*

Monitoring Service Definitions

Purpose Display information about the service definitions configured on your router.

Action To display information for the particular service definition, specify the name of a service definition macro file, including the .mac extension:

```
host1#show service-management service-definition tiered.mac
tiered.mac - WED DEC 14 14:41:20 2005
Installed: True
Service: tiered(inputbw, outputbw)
Reference Count: 0
```

To display summary information for all service definitions:

```
host1#show service-management service-definition brief
Service Definitions
-----
Filename                Service                Installed              Reference
Count
```

```

-----
video.mac      video(inputbw, outputbw)  True      0
tiered.mac     tiered(inputbw, outputbw)  True      0

  Filename           Timestamp
-----
video.mac           TUE NOV 15 15:22:00 2005
tiered.mac          WED DEC 14 14:41:20 2005

```

Meaning Table 176 on page 766 lists the **show service-management service-definition** command output fields.

Table 176: show service-management service-definition Output Fields

Field Name	Field Description
Filename	Name of the service definition macro file
Service	Name of the service, with the parameter specifications in parentheses
Installed	Status of definition: <ul style="list-style-type: none"> • True—installed • False—not installed
Reference Count	Number of times the service definition has been used to instantiate a unique service instance (which identifies the policy, QoS, and profile objects for a service). For example, if one service session—such as, tiered(40000,40000)—is activated by multiple subscribers, the reference count is 1. However, if one subscriber activates tiered(40000,40000) and another subscriber activates tiered(75000,75000)—the reference count is 2.
Timestamp	Day, date, and time the service definition was copied to NVS.

Related Documentation

- *show service-management service-definition*

Monitoring Service Session Profiles

Purpose Display information about service session profiles configured on your router.

Action To display summary information for all service session profiles:

```

host1#show service-management service-session-profile brief
      Service Session Profiles
-----
Name      Volume    Time    Statistics
-----
tiered1   20000    1000    Volume-Time
tiered2   20000    1000    Time

```

```
video1    15000    1000    Volume-Time
video4     0         0      Disabled
```

To display information for a particular service session profile:

```
host1#show service-management service-session-profile tiered1
tiered1
Time       : 1000
Volume     : 20000
Statistics : Time and Volume
```

Meaning Table 177 on page 767 lists the **show service-management service-session-profile** command output fields.

Table 177: show service-management service-session-profile Output Fields

Field Name	Field Description
Name	Name of the service session profile
Volume	Volume threshold, in MB, for the service session
Time	Time threshold, in seconds, for the service session
Statistics	Type of statistics that are captured: <ul style="list-style-type: none"> • Disabled (none) • Time • Volume-Time • Time and Volume

Related Documentation

- *show service-management service-session-profile*

Monitoring Active Owner Sessions with Service Manager

Purpose Display information about active subscriber sessions, by owner.

Action To display summary information for all active owner sessions:

```
host1# show service-management owner-session brief
Subscriber Sessions
```

```
-----
Name           Interface    Id  Owner/Id    State  Non-volatile  Service
-----
CLIENT1@ISP.COM ip192.168.0.3 1  AAA 4194326 Active False      1
CLIENT2@ISP.COM ip192.168.0.7 2  AAA 4194327 Active False      1
CLIENT3@ISP.COM ip192.168.0.4 3  AAA 4194328 Active False      1
CLIENT4@ISP.COM ip192.168.0.5 4  AAA 4194329 Active False      1
CLIENT5@ISP.COM ip192.168.0.6 5  AAA 4194330 Active False      1
CLIENT6@ISP.COM ip192.168.0.8 6  AAA 4194331 Active False      1
CLIENT7@ISP.COM ip192.168.0.1 7  AAA 4194332 Active False      1
CLIENT8@ISP.COM ip192.168.0.9 8  AAA 4194333 Active False      1
```

To display information for a particular owner:

```
host1# show service-management owner-session aaa 4194326
User Name: CLIENT1@ISP.COM, Interface: ip 192.168.0.1
Owner/Id: AAA/4194326
Non-volatile: False
State: Active
ServiceSessions:
      Name                Owner/ID      State                Operation
-----
tiered(2000000,3000000)  AAA 4194326  Config ApplySuccess  Activate
      Name                Non-volatile
-----
tiered(2000000,3000000)  False
```

To display information for a particular owner with service session information:

```
host1# show service-management owner-session aaa 4194326 service-session
User Name: client1@isp.COM, Interface: ip192.168.0.1
Service : tiered(2000000,3000000)
Non-volatile : False
Owner : AAA 4194326
State : Config ApplySuccess
Activate : True
Statistics Type : time-based and volume-based
Statistics Complete : False
Poll Interval : 0
Poll Expire : 0
Activate Time : THU MAR 02 01:21:26 2006
Time : 0
Time Expire : 0
Volume MB: 2
Volume Expire MB : 1
Input Bytes : 594
Output Bytes : 1196
Input Packets : 1
Output Packets : 2
```

Meaning [Table 178 on page 768](#) lists the **show service-management owner-session** command output fields.

Table 178: show service-management owner-session Output Fields

Field Name	Field Description
Name	Name of the subscriber or name of the service session
Interface	Type and IP address of the subscriber's interface
Owner/Id	Method used to activate the subscriber session (CLI, AAA) and ID number generated by the owner
State	Status of the subscriber session (active or inactive), or status of the service session
Non-volatile	Indicates whether the service session is stored in NVS; RADIUS-based service sessions are not stored in NVS

Table 178: show service-management owner-session Output Fields (*continued*)

Field Name	Field Description
Service Sessions	Number of service sessions currently active for this subscriber
Operation	Last operation that Service Manager performed
Service	Name of the service, with parameter values in parentheses
Activate	Indicates whether the last operation was activate (True) or deactivate (False)
Statistics Type	Type of statistics collected; none, time, or volume-time
Statistics Complete	Whether statistics have been successfully collected; True or False
Poll Interval	Interval, in seconds, that interim statistics reports are sent
Poll Expire	Number of seconds until the next statistics report is sent
Activate Time	Day, date, and time when the service session was activated
Time	Time threshold value set by service session profile or RADIUS VSA
Time Expire	Time left until the threshold expires; this value starts as the time threshold value and is decremented as time passes
Volume	Volume threshold value set by service session profile or RADIUS VSA
Volume Expire	Volume left until the threshold is exceeded; this value starts as the volume threshold value and is decremented as the service statistics measure volume
Input Bytes	Current value of input bytes that the statistics configuration is measuring
Output Bytes	Current value of output bytes that the statistics configuration is measuring
Input Packets	Current value of input packets that the statistics configuration is measuring

Table 178: show service-management owner-session Output Fields (*continued*)

Field Name	Field Description
Output Packets	Current value of output packets that the statistics configuration is measuring

Related Documentation • [show service-management owner-session](#)

Monitoring Active Subscriber Sessions with Service Manager

Purpose Display information about active subscriber sessions on your router.

Action To display summary information for all active subscriber sessions:

```
host1# show service-management subscriber-session brief
```

```
Subscriber Sessions
-----
```

Name	Interface	Id	Owner/Id	State	Non-volatile	Service Sessions
CLIENT1@ISP.COM	ip192.168.0.3	1	AAA 4194326	Active	False	1
CLIENT2@ISP.COM	ip192.168.0.7	2	AAA 4194327	Active	False	1
CLIENT3@ISP.COM	ip192.168.0.4	3	AAA 4194328	Active	False	1
CLIENT4@ISP.COM	ip192.168.0.5	4	AAA 4194329	Active	False	1
CLIENT5@ISP.COM	ip192.168.0.6	5	AAA 4194330	Active	False	1
CLIENT6@ISP.COM	ip192.168.0.8	6	AAA 4194331	Active	False	1
CLIENT7@ISP.COM	ip192.168.0.1	7	AAA 4194332	Active	False	1
CLIENT8@ISP.COM	ip192.168.0.9	8	AAA 4194333	Active	False	1
CLIENT9@ISP.COM	ip192.168.0.2	9	AAA 4194334	Active	False	1
CLIENT10@ISP.COM	ip192.168.0.10	10	AAA 4194335	Active	False	1

To display information for that particular subscriber with the subscriber name:

```
host1# show service-management subscriber-session client1@isp.com interface ip 192.168.0.1
```

```
User Name: CLIENT1@ISP.COM, Interface: ip 192.168.0.1
```

```
Id: 1
```

```
Owner: AAA 4194326
```

```
Non-volatile: False
```

```
State: Active
```

```
ServiceSessions:
```

Name	mutex	Owner/Id	State	Operation
tiered(2000000,3000000)		AAA 4194326	ConfigApplySuccess	Activate
Name	Non-volatile			
tiered(2000000,3000000)	False			

To display information for that particular subscriber with the service session:

```
host1# show service-management subscriber-session client1@isp.COM interface ip 192.168.0.1
service-session tiered
```

```
User Name: client1@isp.COM, Interface: ip192.168.0.1
```

```
Service : tiered(2000000,3000000)
```

```
Non-volatile : False
```

```
Owner : AAA 4194326
```

```

State : Config ApplySuccess
Activate : True
Statistics Type : time-based and volume-based
Statistics Complete : False
Poll Interval : 0
Poll Expire : 0
Activate Time : THU MAR 02 01:21:26 2006
Time : 0
Time Expire : 0
Volume MB: 2
Volume Expire MB : 1
Input Bytes : 594
Output Bytes : 1196
Input Packets : 1
Output Packets : 2

```

To display information for a particular subscriber using the subscriber ID:

```
host1#show service-management subscriber-session 20
```

```
User Name: CLIENT50@ISP.COM, Interface: ip192.168.100.33
```

```
Id: 20
```

```
Owner/Id: CLI
```

```
Non-volatile: True
```

```
State: Active
```

```
ServiceSessions:
```

Name	mutex	Owner	State	Operation
internet(5000,8000)	12	CLI	Config ApplySuccess	Activate

Name	Non-volatile
internet(5000,8000)	True

Meaning [Table 179 on page 771](#) lists the **show service-management subscriber-session** command output fields.

Table 179: show service-management subscriber-session Output Fields

Field Name	Field Description
Name	Name of the subscriber or name of the service session
Interface	Type and IP address of the subscriber's interface
Id	ID number of the subscriber session
mutex	Index number of the mutex group to which the service session belongs
Owner/Id	Method used to activate the subscriber session (CLI, AAA) and ID number generated by the owner (Acct-Session-ID for AAA)
State	Status of the subscriber session (active or inactive), or status of the service session
Non-volatile	Indicates whether the service session is stored in NVS; RADIUS-based service sessions are not stored in NVS

Table 179: show service-management subscriber-session Output Fields (*continued*)

Field Name	Field Description
Service Sessions	Number of service sessions currently active for this subscriber
Operation	Last operation that Service Manager performed
Service	Name of the service, with parameter values in parentheses
Activate	Indicates whether the last operation was activate (True) or deactivate (False)
Statistics Type	Type of statistics collected; none, time, or volume-time
Statistics Complete	Whether statistics have been successfully collected; True or False
Poll Interval	Interval, in seconds, that interim statistics reports are sent
Poll Expire	Number of seconds until the next statistics report is sent
Activate Time	Day, date, and time when the service session was activated
Time	Time threshold value set by service session profile or RADIUS VSA
Time Expire	Time left until the threshold expires; this value starts as the time threshold value and is decremented as time passes
Volume	Volume threshold value set by service session profile or RADIUS VSA
Volume Expire	Volume left until the threshold is exceeded; this value starts as the volume threshold value and is decremented as the service statistics measure volume
Input Bytes	Current value of input bytes that the statistics configuration is measuring
Output Bytes	Current value of output bytes that the statistics configuration is measuring
Input Packets	Current value of input packets that the statistics configuration is measuring

Table 179: show service-management subscriber-session Output Fields (continued)

Field Name	Field Description
Output Packets	Current value of output packets that the statistics configuration is measuring

Related Documentation

- *show service-management subscriber-session*

Monitoring the Number of Active Subscriber and Service Sessions with Service Manager

Purpose Display the total number of active subscriber and service sessions configured on your router.

Action To display the total number of active subscriber and service sessions:

```
host1#show service-management summary
```

```
Total Subscriber Sessions : 10
```

```
Total Service Sessions : 10
```

Meaning [Table 180 on page 773](#) lists the **show service-management summary** command output fields.

Table 180: show service-management summary Output Fields

Field Name	Field Description
Total Subscriber Sessions	Number of active subscriber sessions on the router
Total Service Sessions	Number of active service sessions on the router

Related Documentation

- *show service-management summary*

Verifying Computation of Service Session Accounting Based on Scheduler Profiles

Purpose Display whether the capability to compute accounting details based on scheduler profiles for policies, with rate-limit profiles in hierarchical parent groups, on output interfaces is enabled.

Action To determine whether the mechanism to calculate accounting details for subscriber service sessions based on scheduler profiles for policies, with rate-limit profiles in hierarchical parent groups, on output interfaces is enabled:

```
host1#show service-accounting-statistics
```

```
Service accounting mode is scheduler based.
```

**Related
Documentation**

- [Configuring Calculation of Service Session Accounting Based on Scheduler Profiles Instead of Rate-Limit Profiles in Hierarchical Parent Groups for Forwarded Packets on page 703](#)
- *service-accounting-statistics scheduler-based*
- *show service-accounting-statistics*

PART 7

Index

- [Index on page 777](#)

Index

Symbols

10-Gigabit Ethernet interfaces	
specifying an interface.....	653

A

AAA (authentication, authorization, accounting)	
DSL Forum VSAs.....	206, 217
EAP authentication.....	17
failure, notifying RADIUS of.....	39
L2TP tunnel switch profiles,	
applying.....	392, 393
overview.....	5, 305
services	
accounting.....	305
and TACACS+.....	305
authentication.....	305
authorization.....	305
overview.....	306
AAA authentication	
scenarios in which show subscribers command	
output varies	
domain name mapped with the virtual	
router.....	482
domain name not mapped with the virtual	
router.....	482
aaa commands.....	22
aaa accounting acct-stop on-aaa-failure.....	39
aaa accounting acct-stop	
on-access-deny.....	39
aaa accounting broadcast.....	22
aaa accounting commands.....	310
aaa accounting default.....	22
aaa accounting duplication.....	22
aaa accounting exec.....	310
aaa accounting immediate-update.....	22
aaa accounting interim-update.....	21, 89
aaa accounting interval.....	22
aaa accounting statistics.....	22
aaa accounting suppress null-username.....	310
aaa accounting vr-group.....	22
aaa authentication default.....	74

aaa authentication enable default.....	305, 310
aaa delimiter.....	9
aaa dhcpv6-delegated-prefix	
delegated-ipv6-prefix.....	39
aaa dns primary.....	24
aaa dns secondary.....	24
aaa domain-map.....	7, 8, 16, 348
aaa intf-desc-format include.....	211
aaa ipv6-nd-ra-prefix framed-ipv6-prefix.....	39
aaa local database.....	74
aaa local select database.....	74
aaa local username	74
aaa new-model.....	305
aaa parse-direction.....	9
aaa parse-order.....	9
aaa profile.....	29, 31
aaa radius-override-ncp-negotiation.....	107
aaa route-download.....	29
aaa route-download now.....	29
aaa route-download suspend.....	29
aaa service accounting interval.....	700
aaa subscriber limit per-port.....	39
aaa subscriber limit per-vr.....	39
aaa timeout.....	38
aaa tunnel assignment-id-format.....	348
aaa tunnel calling-number-format.....	338
aaa tunnel calling-number-format	
fallback.....	338
aaa tunnel client-name.....	348
aaa tunnel ignore.....	348
aaa tunnel password.....	348
aaa tunnel switch-profile.....	394
aaa tunnel tx-connect-speed-method	401
aaa tunnel-group.....	352
aaa user accounting interval.....	21, 89, 700
aaa virtual-router.....	89
aaa wins primary.....	24
aaa wins secondary.....	24
See also show aaa commands	
AAA default tunnel parameters	
L2TP transmit connect speed.....	400
AAA domain maps	
L2TP transmit connect speed.....	399
preference order	
for determining local address pools.....	54
tunnel subscribers.....	23
AAA LLID (logical line identifier).....	31
configuration steps.....	31
how it works.....	31

monitoring.....	124, 132
preauthentication considerations.....	31
RADIUS attributes in preauthentication request.....	31, 33
troubleshooting.....	31
using to track subscribers.....	31
AAA logical line identifier (LLID). See AAA LLID	
AAA profile commands	
allow.....	29
deny.....	29
ppp aaa-profile.....	31
AAA profiles.....	29
allowing or denying domain names.....	29
configuring.....	29
creating domain name aliases.....	29
AAA tunnel groups	
L2TP transmit connect speed.....	399
Access-Accept messages	182
preference order	
in allocation of IPv6 prefixes.....	55
with Framed-IPv6-Prefix attribute for Prefix Delegation.....	39
with IPv6-NdRa-Prefix attribute for IPv6 Neighbor Discovery.....	39
Access-Challenge messages.....	182, 208
Access-Reject messages.....	182, 208
Access-Request messages.....	182
ANCP (L2C)-related VSAs.....	215
DSL Forum VSAs.....	206, 217
accounting	
broadcast.....	17
configuring servers.....	17
configuring TACACS+.....	305
description.....	5
duplicate.....	17
interim.....	21
server access.....	17
server request processing limit.....	17
specifying methods.....	17
TACACS+.....	305
accounting statistics	
tunneled PPP session.....	403
Acct-Authentic (RADIUS attribute 45).....	210
Acct-Delay-Time (RADIUS attribute 41).....	210
Acct-Input-Gigapackets (RADIUS attribute 26-35).....	213
Acct-Input-Gigawords (RADIUS attribute 52).....	211
Acct-Link-Count (RADIUS attribute 51).....	211
Acct-Multi-Session-Id (RADIUS attribute 50).....	210
Acct-Off messages.....	193
Acct-On messages.....	193
Acct-Output-Gigapackets (RADIUS attribute 26-36).....	213
Acct-Session-Id (RADIUS attribute 44).....	210, 688, 698
Acct-Start messages.....	193
ANCP (L2C)-related VSAs.....	215
DSL Forum VSAs.....	206, 217
Acct-Stop messages.....	193
ANCP (L2C)-related VSAs.....	215
DSL Forum VSAs.....	206, 217
Acct-Terminate-Cause (RADIUS attribute 49).....	210
Acct-Tunnel-Connection (RADIUS attribute 68).....	211
Activate-Service (RADIUS attribute 26-65).....	687
active subscriber details, viewing	
with AAA authentication configured on default router	
show subscribers command.....	482
with RADIUS authentication configured on default router	
show subscribers command.....	482
show subscribers ipv6 command.....	482
show subscribers summary command.....	482
address command, L2TP.....	348, 352
address pool	
ranges.....	26
agent circuit ID (suboption 1).....	526
agent remote ID (suboption 2).....	526
agent-circuit-id	
including in Calling Number AVP.....	338
including in PPPoE remote circuit ID format.....	520
agent-remote-id	
including in Calling Number AVP.....	338
including in PPPoE remote circuit ID format.....	520
allow command.....	29
Ascend-Num-In-Multilink (RADIUS attribute 188).....	212
ATM (Asynchronous Transfer Mode)	
E120 and E320 routers.....	622
atm atm1483 advisory-rx-speed command and L2TP.....	355
atm atm1483 commands	
atm atm1483 advisory-rx-speed.....	356

- atm commands
 - atm.....34
 - atm pvc653
- ATM subinterface
 - configuring multiple clients.....84
 - configuring single clients.....83
- attribute value pair. *See* AVP
- authentication
 - AAA overview.....305
 - configuring servers.....17
 - configuring TACACS+.....305
 - description.....5
 - EAP.....17
 - preauthenticating users.....8
 - redirected authentication.....8
 - server access.....17
 - server request processing limit.....17
 - specifying methods.....17
- authentication and accounting servers
 - configuring.....17
- authentication login, TACACS+.....305
- authentication of IPv6 clients
 - and behavior of the show subscribers
 - command.....482
- authentication, authorization, accounting. *See* AAA
- authorization
 - AAA overview.....305
 - description.....5
 - TACACS+.....305
- AVP (attribute value pair).....322
 - Bearer Type (AVP 18)
 - relaying in L2TP tunnel-switched
 - network.....389
 - relaying in L2TP
 - tunnel-switchednetwork.....391
 - Calling Number (AVP 22)
 - formatting and preventing in ICRQ
 - packets.....338
 - relaying in L2TP tunnel-switched
 - network.....389
 - relaying in L2TP
 - tunnel-switchednetwork.....391
 - Cisco NAS Port Info (AVP 100)
 - relaying in L2TP tunnel-switched
 - network.....389
 - relaying in L2TP
 - tunnel-switchednetwork.....391
 - Transmit (TX) Speed (AVP 24)
 - reporting transmit connect speed in.....394
- avp command.....394
- B**
- B-RAS applications
 - AAA profiles.....29
 - allowing or denying domain names.....29
 - client to server interaction.....17
 - configuring
 - authentication and accounting servers.....17
 - B-RAS license.....67
 - DHCP external server application.....545
 - DHCP features.....461
 - IP addresses for remote clients.....4
 - local address servers.....25
 - name server addresses.....24
 - SRC client.....47
 - timeout.....38
 - UDP checksums.....22
 - creating an IP interface.....83
 - creating domain name aliases.....29
 - DHCP (Dynamic Host Configuration Protocol)
 - proxy client and server.....4
 - IP hinting.....8
 - limiting active subscribers.....39
 - local address server.....4
 - mapping user domain names to a virtual
 - router.....7
 - mapping user requests
 - without a valid domain name.....7
 - without configured domain name.....8
 - monitoring.....111
 - multiple clients per ATM subinterface.....84
 - overview.....4
 - preauthenticating users.....8
 - protocol support.....6
 - redirected authentication.....8
 - references.....181, 262
 - single clients per ATM subinterface.....83
 - specifying a single name for a domain.....16
 - SRC client. *See* SRC software
 - virtual router.....7
- B-RAS licenses
 - configuring.....67
- baseline commands
 - baseline aaa.....113
 - baseline aaa route-download.....113
 - baseline cops.....113
 - baseline dhcp relay.....560
 - baseline dhcp server.....560

baseline dhcp-local.....	561
baseline dhcpv6-local.....	561
baseline ip dhcp-external.....	561
baseline ip http.....	749
baseline local pool.....	113
baseline radius.....	113
baseline radius dynamic-request.....	298
baseline radius relay.....	301
baseline ssc.....	113
Bearer Type AVP	
relaying in L2TP tunnel-switched	
network.....	389, 391
BOOTP (bootstrap protocol).....	512
bootstrap protocol. See BOOTP	
bridged Ethernet and dynamic subscriber	
interfaces.....	647
Broadband Remote Access Server. See B-RAS	
applications	
broadcast AAA accounting.....	17
configuring.....	17
broadcast flag, DHCP	
controlling transmission of DHCP reply	
packets.....	516
interaction with layer 2 unicast transmission	
method.....	516
bundled session commands	
bundled-group-id.....	366, 377
bundled-group-id-overrides-mlppp-ed.....	366, 377
bundled sessions.....	377
C	
cable modem DHCP relay.....	525
cable modem networks.....	626
Called-Station-Id (RADIUS attribute 30).....	210
Calling Number AVP	
descriptive formats.....	338
fixed format.....	338
fixed format configuration.....	338
formatting in L2TP ICRQ packets.....	338
including agent-circuit-id and	
agent-remote-id.....	338
preventing in L2TP ICRQ packets.....	338
relaying in L2TP tunnel-switched	
network.....	389, 391
Calling-Station-Id (RADIUS attribute 31).....	210
captive portal. See guided entrance	
Change-of-Authorization-Request messages.....	182
Cisco NAS Port Info AVP	
relaying in L2TP tunnel-switched	
network.....	389, 391
Class (RADIUS attribute 25).....	210
clear ip commands	
clear ip dhcp-local binding	491
clear ip routes download.....	29
clear ipv6 routes download.....	29
CLI (command-line interface)	
authorization and authentication	
messages.....	208
commands used to modify RADIUS	
attributes.....	209
client-name command.....	351, 354
COA-Request messages.....	182
guided entrance.....	731
Service Manager.....	731
combined IPv4 and IPv6 services	
example	
with input traffic flow.....	740
with output traffic flow.....	740
example scenario	
for rate limiting VoIP traffic.....	740
external parent groups and	
example.....	740
in a dual stack	
activating.....	696
backward compatibility.....	696
deactivating.....	696
example.....	740
performance impact.....	696
rate limiting and	
example.....	740
service interim accounting.....	705
statistics collection and	
external parent groups.....	722
command-line interface. See CLI	
Common Open Policy Service. See COPS	
compatibility with previous releases	
for IPv4 and IPv6 services	
in a dual stack, combined.....	697
in a dual stack, independent.....	697
configuring. See specific feature, product, or	
protocol	
Connect-Info (RADIUS attribute 77).....	211
conventions	
notice icons.....	xxxiii
text and syntax.....	xxxiv

COPS (Common Open Policy Service).....	47, 138, 140
COPS(Common Open Policy Service).....	138, 140
COPS-PR (COPS usage for policy provisioning).....	47
customer support.....	xxxvi
contacting JTAC.....	xxxvi

D

Deactivate-Service (RADIUS attribute 26-66).....	687, 691, 735
default domain name.....	7
default-router command.....	653
default-upper-type mlppp command.....	369
Delegated-Ipv6-Prefix (RADIUS attribute 123).....	212
delegating routers	
allocation of IPv6 prefixes to requesting routers.....	51
as E Series routers for allocation of prefixes.....	51
assigning prefixes to DHCPv6 clients.....	101
to ndra clients.....	105
deny command.....	29
descriptive formats	
Calling Number AVP.....	338
destination	
changing.....	323
DHCP (Dynamic Host Configuration Protocol)	
access model.....	464
configuring BOOTP Relay.....	512
configuring DHCP relay.....	511
configuring DHCP relay proxy.....	541, 542
configuring proxy client and server.....	464
features.....	28
giaddr.....	473, 476, 512
option	
82,497,498,499,512,514,515,520,529,530,531,533,536	
overview.....	28, 461, 511
per-interface logging.....	465
source IP address.....	512
trust-all.....	512
DHCP access model	
configuring.....	461
DHCP broadcast flag	
interaction with layer 2 unicast transmission method.....	516
DHCP client bindings	
deleting.....	466
indexing using DUID.....	478
managing.....	466
requests from clients	
acknowledged using default DUID type by server.....	478
viewing.....	466
DHCP clients	
maximum number of prefixes allocated to using DHCPv6 local server.....	40
DHCP commands	
dhcp-external	
duplicate-mac-address.....	548, 557
ip auto-detect ip-subscriber.....	653
ip dhcp-capture.....	465
ip dhcp-external	
re-authenticate-subscriber-interface.....	558
ip dhcp-external	
recreate-subscriber-interface.....	547, 553
ip inactivity-timer.....	653
show dhcp-external.....	574
dhcp delete-binding command.....	468
DHCP external server	
and DHCP relay proxy.....	546
configuring.....	546, 550, 609
dynamic subscriber interfaces	
interoperating with DHCP relay and DHCP relay proxy	554
identifying clients with duplicate MAC addresses.....	548, 557
monitoring.....	570
preserving dynamic subscriber interfaces.....	547, 553
re-authenticating auto-detected dynamic subscriber interfaces.....	558
DHCP local address pools	
configuring.....	491
grace periods.....	494
linking.....	493
DHCP local server.....	488
configuring.....	485
configuring authentication.....	497
duplicate clients.....	488
equal-access mode.....	472
address allocation.....	472
configuring.....	477
connection process.....	472
local pool selection.....	472

overview.....	472	DHCP per-interface information	
SRC (Session and Resource Control		logging.....	465
software).....	461	DHCP pool commands	
local address pool group.....	493, 577	default-router.....	494, 653
local pool selection, equal-access.....	472	dns-server.....	494
using domain name.....	473	domain-name.....	494
using framed IP address.....	473	grace-period.....	494
using giaddr.....	473	ip dhcp-local pool.....	653
using pool name.....	473	lease.....	494
local pool selection, standalone.....	474	link.....	494
using giaddr.....	474	netbios-name-server.....	494
using received interface IP address.....	474	netbios-node-type.....	494
logging information.....	491	network.....	494
modes.....	471	reserve.....	494
monitoring.....	571	server-address.....	494
overview.....	471	snmpTrap.....	494
RADIUS accounting support for.....	471	use-release-grace-period.....	494
RADIUS accounting support for standalone		warning.....	494
mode.....	45	DHCP proxy client	
standalone mode		configuring.....	464
address allocation.....	474, 497	DHCP relay	
authentication.....	497	configuring.....	512
configuring.....	477	creating.....	512
local pool selection.....	474	in same VR as dynamic subscriber	
overview.....	474	interfaces.....	628
unique client IDs.....	488	interoperating with DHCP external server.....	554
DHCP local server traps		logging information.....	521
logging.....	490	preventing host route installation.....	518
DHCP local server:local pool selection, equal-access		removing.....	512
using received interface IP address.....	473	spoofed giaddr.....	514
DHCP local server:local pool selection, standalone		spoofed relay agent option.....	514
with AAA		DHCP relay agent information option	
using domain name.....	476	agent circuit ID (suboption 1).....	530
using framed IP address.....	475	agent remote ID (suboption 2)	530
using giaddr.....	476	vendor-specific (suboption 9).....	530
using pool name.....	476	DHCP relay and BOOTP relay	
using received IP address.....	476	configuring.....	512
DHCP local server:local pool selection, standalone		DHCP relay proxy	
without AAA		best offer.....	542
using giaddr.....	475	configuring.....	541
using received interface IP address.....	475	first offer.....	542
DHCP logging.....	584	interoperating with DHCP external server.....	554
DHCP option 60 cable modem DHCP relay.....	525	DHCP server	
DHCP option 60 strings.....	522	dynamic subscriber interfaces.....	628
DHCP Option 82 (RADIUS attribute 26-159).....	215	DHCP Subscribers	
DHCP options and RADIUS		Sending RADIUS Attributes.....	495
configuring.....	461	DHCP unique identifier See DUID	
DHCP packets		DHCP vendor class identifier option.....	522
logging.....	584		

dhcp-external commands	
dhcp-external	
duplicate-mac-address.....	548, 557
show dhcp-external.....	574
DHCP-GI-Address (RADIUS attribute 26-57).....	214
dhcp-local-server-configuration-overview.....	477
DHCP-MAC-Address (RADIUS attribute 26-56).....	214
DHCP-Options (RADIUS attribute 26-55).....	214
DHCPv6 client bindings	
deleting.....	505
DHCPv6 clients	
assigning prefixes to	
using local address pools.....	101
dhcpv6 delete-binding command.....	506
DHCPv6 local address pools See IPv6 local address pools	
DHCPv6 local server	
accounting for IPv6 subscribers	
standalone mode with AAA authentication.....	481
Advertise message	
authentication of subscribers.....	481
assigning maximum number of IPv6 prefixes	
using Prefix Delegation.....	40
assigning prefixes to clients in conjunction with Neighbor Discovery.....	41
authentication of IPv6 subscribers	
connected using Ethernet links.....	480
connected using VLAN links.....	480
configuring.....	501
configuring authentication.....	500
configuring DUID type.....	504
IPv6.....	501
limiting maximum number of IPv6 prefixes to be delegated.....	483
maximum number of IPv6 prefixes assigned to clients, using Prefix Delegation and Neighbor Discovery.....	40
to clients, using Prefix Delegation only.....	40
prefix delegation	
with standalone mode configured.....	480
with standalone mode for PPP subscribers.....	480
Sending Delegated-IPv6-Prefix to SRC client	
overview.....	56
verifying.....	602
standalone mode	
address allocation.....	500
authentication.....	500
standalone mode with AAA authentication	
construction of username.....	480
standalone mode with AAA authentication, overview.....	480
standalone mode without AAA authentication, overview.....	480
using DUID in identity verification of clients.....	478
using DUID types that suit the client and service provider networks.....	478
DHCPv6 Prefix Delegation	
and IPv6 Neighbor Discovery	
without configuring	
Delegated-IPv6-Prefix.....	39
and Neighbor Discovery for prefixes delegation	
scaling limit, same prefix for multiple subscribers.....	41
scaling limit, unique prefix per subscriber.....	41
assigned prefix length of /128	
in local address pools.....	53
enabling	
IPv6 local address pool feature.....	101
example for non-PPP client requests.....	104
example scenario.....	54
for client requests	
over non-PPP links.....	52
over PPP links.....	52
for IPv6 clients	
overview.....	51
Framed-IPv6-Prefix	
in Access-Accept messages.....	39
guidelines for configuring	
IPv6 local address pools.....	52
interface level configuration	
versus router level configuration.....	55
limitation on	
prefixes assigned to clients.....	103
maximum number of prefixes delegated to clients.....	40
See also using DHCPv6 local server only	

standard RADIUS attributes	
configuring	39
verifying.....	154
using IPv6 local address pools	
monitoring a single pool.....	170
monitoring all configured pools.....	169
monitoring statistics for a pool.....	171
dialed number identification service. See DNS	
digital subscriber line access multiplexers. See DSLAMs	
digital subscriber lines. See DSLs	
disable proxy lcp command.....	369
Disconnect-Cause (RADIUS attribute 26-51).....	213
Disconnect-Request messages.....	182
DNIS (dialed number identification service).....	8, 368
DNS (Domain Name System)	
assigning IP addresses.....	123
overview.....	24
DNS addresses	
order of preference	
in allocation to clients.....	55
DNS domain names	
list of	
configured in IPv6 local address pools.....	55
order of preference	
in responses to clients.....	55
DNS domains	
configuring more than one	
using the CLI interface.....	102
in IPv6 local address pools	
processing client requests for resolution.....	102
in responses to clients	
Domain Search List option and.....	102
maximum number	
in IPv6 local address pools.....	102
DNS Recursive Name Search option	
DHCPv6 server responses	
and DNS servers in local pools.....	102
DNS servers	
addresses in responses to clients	
DNS Recursive Name Search option	
and.....	102
configuring in	
IPv6 local address pools.....	102
list of	
configured in IPv6 address pools.....	55
order of preference	
in responses to clients.....	55
order of use	
for delegating prefixes.....	102
primary and secondary	
for domain resolution requests from clients.....	102
responding with IPv6 addresses	
for client requests.....	102
documentation set	
comments on.....	xxxv
Domain Name System. See DNS	
domain names	
allowing or denying.....	29
configuring.....	9
default.....	7
mapping to virtual routers.....	7, 120, 131, 427
mapping user requests without domain name.....	7, 8
none.....	8
specifying single name for users.....	16
stripping domain name.....	9
using aliases.....	29
using delimiters other than @.....	9
using either domain or realm as domain name.....	9
using realm name as domain name.....	9
Downstream-Calculated-QoS-Rate (RADIUS attribute 26-141).....	214
DS-Lite-Tunnel-Name (RADIUS attribute 144).....	212
DSL Forum VSAs	
controlling inclusion of.....	217
descriptions.....	260
in AAA access and accounting messages.....	206
DSLAMs (digital subscriber line access multiplexers).....	4
DSLs (digital subscriber lines).....	5
dual stack	
combined IPv4 and IPv6 services	
example of	740
IPv4 and IPv6 services	
combined, activating and deactivating.....	696
combined, overview.....	694
independent, activating and deactivating.....	696
independent, overview.....	694
service interim accounting, overview.....	705

Service manager support, activating and deactivating.....	696
Service manager support, overview.....	694
statistics collection	
external parent groups.....	722
dual-stack	
IPv6 service accounting statistics	
separate collection.....	706
DUID	
conformance standards.....	478
default type on DHCPv6 local server.....	478
format and example.....	478
Identity Association for Prefix Delegation option and.....	478
Identity Association identifier and.....	478
indexing client bindings.....	478
prerequisite for configuring DUID type.....	504
type, configuring on DHCPv6 local server.....	504
types of	
Type 1, contents and format.....	478
Type 2, contents and format.....	478
Type 3, contents and format.....	478
types supported	
by DHCP client.....	478
by DHCPv6 local server.....	478
used in client and server exchanges.....	478
duplicate AAA accounting.....	17
configuring.....	17
duplicate clients.....	488
duplicate MAC addresses	
identifying DHCP clients with.....	548, 557
monitoring.....	574
Dynamic Host Configuration Protocol. <i>See</i> DHCP	
<i>See</i> DHCP	
dynamic IP interfaces.....	34
dynamic subscriber interfaces	
commands.....	653
configuring.....	628, 642
configuring DHCP external server to interoperate with DHCP relay and DHCP relay proxy	554
configuring DHCP external server to preserve.....	547, 553
configuring DHCP external server to re-authenticate.....	558
DHCP server.....	628
framed routes.....	632
GRE tunnel configuration.....	647
in same VR as DHCP relay.....	628

inheriting MAC validation state.....	632
IP over bridged Ethernet configuration.....	645
IP over Ethernet configuration.....	642
IP over VLAN over Ethernet configuration.....	643
monitoring.....	655
overview.....	628
packet detection.....	631
E	
E120 and E320 routers	
ATM interfaces.....	622
EAP (Extensible Authentication Protocol)	
external RADIUS server.....	17
local authentication server.....	17
RADIUS attributes.....	17
RADIUS authentication.....	17
TACACS+ server.....	17
EAP-Message (RADIUS attribute 79).....	20
Egress-Policy-Name (RADIUS attribute 26-11).....	213
enable proxy authenticate command.....	369
encapsulation commands	
encapsulation bridge1483.....	653
encapsulation vlan.....	653
endpoint discriminator.....	377
equal-access DHCP local server.....	472
Ethernet	
configuring dynamic subscriber	
interfaces.....	642
Ethernet interfaces	
commands	
interface tenGigabitEthernet.....	653
Ethernet links	
between CPE and PE routers	
pool section for Prefix Delegation.....	104
Event-Timestamp (RADIUS attribute 55).....	211
exclusion ranges	
configuring	
for delegation of prefixes.....	102, 106
example for non-PPP client requests.....	104
for DHCPv6 prefixes	
delegated to clients.....	102
for ndra prefixes	
delegated to clients.....	106
Extensible Authentication Protocol. <i>See</i> EAP	
external parent groups	
combined IPv4 and IPv6 services with	
example.....	740
statistics collection for	
setting up.....	722

F

fixed format	
Calling Number AVP.....	338
fragmentation	
and reassembly.....	325
packet.....	323
framed routes dynamic subscriber interfaces.....	632
Framed-Compression (RADIUS attribute 13).....	210
Framed-Interface-Id (RADIUS attribute 96).....	211
Framed-Ip-Address (RADIUS attribute 8).....	210
Framed-Ip-Netmask (RADIUS attribute 9).....	210
Framed-Ipv6-Pool (RADIUS attribute 100).....	211
Framed-Ipv6-Prefix (RADIUS attribute 97).....	211
Framed-IPv6-Prefix attribute	
configuring the same IPv6 prefix for multiple subscribers	
in the Access-Accept message.....	41
used for DHCPv6 Prefix Delegation	
from Access-Accept messages.....	39
Framed-Ipv6-Route (RADIUS attribute 99).....	211
Framed-MTU (RADIUS attribute 12).....	20
Framed-Route (RADIUS attribute 22).....	210

G

giaddr.....	473, 512
GRE (Generic Routing Encapsulation) tunnels	
dynamic subscriber interfaces.....	631, 648
guided entrance.....	663, 730
COA-Request messages.....	731

H

HTTP local server.....	731, 735
guided entrance service.....	735
Service Manager.....	735

I

ICR partition accounting	
viewing the status, enabled or disabled.....	304
identification command.....	351, 354
Identity Association for Prefix Delegation	
and its interoperation with DUID.....	478
contains Identity Association identifier	
collection of prefixes.....	478
Identity Association identifier	
chosen by the requesting router.....	478
using to denote Identity Association for Prefix Delegation options.....	478
idle timeout for B-RAS	
configuring.....	38

idle timeout, range for.....	38
independent IPv4 and IPv6 services	
in a dual stack	
activating.....	696
backward compatibility.....	696
deactivating.....	696
performance impact.....	696
service interim accounting.....	705
statistics collection and	
external parent groups.....	722
Ingress-Policy-Name (RADIUS attribute 26-10).....	213
inheriting MAC address validation state.....	632
interface commands	
interface atm.....	653
interface fastEthernet.....	653
interface gigabitEthernet.....	653
interface ip.....	639
interface loopback.....	653
interface tenGigabitEthernet.....	653
Interface-Desc (RADIUS attribute 26-63).....	214
interfaces	
configuring for DHCP local server.....	477
moving.....	625
Interim-Acct messages.....	193
ANCP (L2C)-related VSAs.....	215
DSL Forum VSAs.....	206, 217
Internet Protocol. <i>See</i> IP	
IOAs	
including in RADIUS Calling-Station-Id	
format.....	210
IP	
hinting.....	8
IP addresses	
assigning to name servers.....	24, 123
configuring for remote client.....	4
ip commands	
ip address.....	653
ip address-pool dhcp.....	464
ip auto-configure	
append-virtual-router-name.....	653
ip auto-configure ip-subscriber.....	630, 653
ip demux-type.....	639
ip destination-prefix.....	638, 639
ip dhcp-local pool.....	653
ip share-interface.....	638
ip share-nexthop.....	638
ip source-prefix.....	640, 653
ip unnumbered.....	653

- ip use-framed-routes ip-subscriber.....653
- ip-hint.....8
- ip dhcp-external commands.....552
 - ip dhcp-external auto-configure.....552
 - ip dhcp-external
 - disregard-giaddr-next-hop.....551
 - ip dhcp-external
 - recreate-subscriber-interface.....547, 553
 - ip dhcp-external server-address.....551
 - ip dhcp-external server-sync.....551
 - ip re-authenticate-auto-detect
 - ip-subscriber.....558
 - See also* show ip dhcp-external commands
- ip dhcp-local auth domain command.....499, 504
- ip dhcp-local auth include command.....499
- ip dhcp-local auth password.....499
- ip dhcp-local auth user-prefix command.....499
- ip dhcp-local commands
 - ip dhcp-local auto-configure
 - agent-circuit-identifier.....491
 - ip dhcp-local excluded-address.....491
 - ip dhcp-local limit.....491
 - ip dhcp-local pool.....494
 - ip dhcp-local unique-client-ids.....491
- ip dhcp-server commands
 - ip dhcp-server.....464
- ip http commands
 - ip http.....735
 - ip http access-class.....735
 - ip http max-connection-time.....735
 - ip http port.....735
 - ip http redirecturl.....735
 - ip http same-host-limit.....735
 - ip http server.....735
- IP interfaces
 - creating.....639
- IP interfaces that support PPP clients
 - configuring.....83
- IP spoofing
 - preventing.....632
- IPCP negotiation
 - based on RADIUS returned attributes.....60
 - based on RADIUS returned attributes and
 - subscriber domain map configuration.....60
- IPv4 and IPv6 services.....697
 - combined services in a dual stack
 - example.....740
 - in a dual stack
 - and Service Manager support.....697
 - See also* combined IPv4 and IPv6 services
 - See also* independent IPv4 and IPv6 services
- IPv4 services
 - in a dual stack
 - activating.....696
 - combined and independent
 - configuration.....694
 - deactivating.....696
 - with IPv6 services.....694
- ipv6 commands
 - ipv6 virtual-router.....8
 - ipv6-local-interface.....8
- IPv6 DHCP local server.....501
 - monitoring.....593, 595, 600
- ipv6 dhcpv6-local auth domain command.....501
- ipv6 dhcpv6-local auth password.....501
- ipv6 dhcpv6-local auth user-prefix command.....501
- ipv6 dhcpv6-local commands
 - ipv6 dhcpv6-local delegated-prefix.....501
 - ipv6 dhcpv6-local dns-domain-search.....501
 - ipv6 dhcpv6-local dns-server.....501
 - ipv6 dhcpv6-local limit.....501
 - ipv6 dhcpv6-local prefix-lifetime.....501
- IPv6 local address pools
 - assigned prefix length of /128
 - Prefix Delegation and.....53
 - configuring
 - for Ndra.....105
 - for Prefix Delegation.....100
 - DNS servers in
 - to return to clients.....102
 - enabling.....101, 105
 - example for non-PPP client requests.....104
 - for delegation of prefixes
 - overview.....51
 - for DHCPv6 Prefix Delegation
 - single pool details, viewing.....170
 - statistics for a single pool, viewing.....171
 - summary of all configured pools,
 - viewing.....169
 - for NDRA
 - single pool details, viewing.....173
 - summary of all configured pools,
 - viewing.....172, 174

guidelines for configuration.....	52	IPv6 prefix ranges	
limitation on		configuring	
number of allocated prefixes.....	103	with the starting and ending prefixes.....	101
multiple configuration		with the starting prefix and length.....	101
on virtual router, preference order.....	54	IPv6 prefix ranges for ndra	
not configured in domain map		configuring	
method for determining prefix to be		with the starting and ending prefixes.....	106
delegated.....	54	with the starting prefix and length.....	106
order of preference		IPv6 prefixes	
in selection for delegation of prefixes.....	54	common prefix for multiple subscribers	
Prefix Delegation		assigned	
example scenario.....	54	using DHCPv6 local server and Neighbor	
prerequisite for configuring.....	101, 105	Discovery.....	41
procedure		maximum number assigned to clients	
for configuring on a virtual		using DHCPv6 local server only.....	41
router.....	100, 105	same prefix with multiple next-hops	
specifying		assigned to IPv6 clients.....	41
domain name for DNS		unique prefix per subscriber assigned	
resolution.....	100, 105	using DHCPv6 local server and Neighbor	
exclusion range for prefixes.....	100, 105	Discovery.....	41
IPv6 address of DNS server.....	100, 105	IPv6 service accounting statistics separate	
preferred lifetime.....	100	collection	
prefix range.....	100, 105	in a dual stack.....	706
starting and ending prefixes of a		IPv6 services	
range.....	100, 105	in a dual stack	
valid lifetime.....	100	activating.....	696
used for Prefix Delegation from		combined and independent	
AAA domain map.....	54	configuration.....	694
interface address.....	54	deactivating.....	696
RADIUS server.....	54	with IPv4 services.....	694
IPv6 Neighbor Discovery		IPv6 subscribers	
and DHCPv6 Prefix Delegation		accounting of	
without configuring Delegated-IPv6-Prefix		using DHCPv6 local server, standalone	
.....	39	mode.....	480
assigning prefixes to clients		accounting of, DHCPv6 local server	
maximum number permissible, same prefix		configured with standalone authentication	
for multiple clients.....	41	mode.....	481
maximum number permissible, unique		authentication of	
prefix per client.....	41	using DHCPv6 local server, standalone	
IPv6-NdRa-Prefix		mode.....	480
in Access-Accept messages.....	39	IPv6-Acct-Input-Gigawords [26-155]	215
maximum number of delegated IPv6 prefixes		IPv6-Acct-Input-Octets [26-151]	215
for requesting clients.....	40	IPv6-Acct-Input-Packets [26-153].....	215
standard RADIUS attributes		IPv6-Acct-Output-Gigawords [26-156].....	215
configuring	39	IPv6-Acct-Output-Octets [26-152].....	215
verifying.....	154	IPv6-Acct-Output-Packets [26-154]	215
		IPv6-Egress-Policy-Name (RADIUS attribute	
		26-107).....	214

- Ipv6-Ingress-Policy-Name (RADIUS attribute 26-106).....214
- Ipv6-Local-Interface (RADIUS attribute 26-46).....213
- Ipv6-NdRa-Prefix (RADIUS attribute 26-46).....214
- IPv6-NdRa-Prefix attribute
 - used for IPv6 Neighbor Discovery
 - from Access-Accept messages.....39
- IPv6-Primary-DNS (RADIUS attribute 26-47).....213
- Ipv6-Secondary-DNS (RADIUS attribute 26-46).....213
- IPv6-Virtual-Router (RADIUS attribute 26-45).....213
- IPv6CP negotiation
 - based on RADIUS returned attributes.....60
 - based on RADIUS returned attributes and subscriber domain map configuration.....60
- L**
- L2C-Down-Stream-Data (RADIUS attribute 26-93).....214
- L2C-Information (RADIUS attribute 26-81).....214
- L2C-Up-Stream-Data (RADIUS attribute 26-92).....214
- L2TP (Layer 2 Tunneling Protocol)
 - actual upstream rate transmission.....354
 - advisory receive speed transmission.....354
 - defining.....321
 - hello packets fast path processing.....402
 - high availability considerations.....329
 - implementation.....323
 - license.....328
 - modifying LAC default settings.....332, 380
 - monitoring.....427
 - peer resynchronization.....385
 - RX Connect-Speed AVP generation.....354
 - rx speed355
 - sessions supported.....328
 - silent failover.....385
 - tunnel selection.....361
 - tunnel switch profiles.....389
- L2TP access concentrator. *See* LAC
- l2tp commands.....381
 - disconnect-cause.....381
 - failover-resync389
 - l2tp checksum.....333
 - l2tp destination lockout-test.....359
 - l2tp destination lockout-timeout.....359
 - l2tp destination profile.....366, 370
 - l2tp destruct-timeout.....334
 - l2tp disable calling-number avp.....338
 - l2tp disable challenge.....348
 - l2tp disable tunnel-hello.....337
 - l2tp disconnect-cause.....381
 - l2tp drain.....334
 - l2tp drain destination.....334
 - l2tp drain tunnel.....334
 - l2tp fail-over-within-preference.....361
 - l2tp failover-resync385, 389
 - l2tp fastpath-hello-processing.....402
 - l2tp ignore-receive-data-sequencing.....348
 - l2tp ignore-transmit-address-change.....360
 - l2tp reject-transmit-address-change.....360
 - l2tp retransmission.....337
 - l2tp rx-connect-speed-upstream-rate356
 - l2tp rx-connect-speed-when-equal356
 - l2tp short-drain-timeout.....334
 - l2tp shutdown.....336
 - l2tp shutdown destination.....336
 - l2tp shutdown session.....336
 - l2tp shutdown tunnel.....336
 - l2tp switch-profile.....394
 - l2tp tunnel default-receive-window.....382
 - l2tp tunnel idle-timeout.....379
 - l2tp tunnel test.....380
 - l2tp tunnel-switching.....379, 394
 - l2tp unlock destination.....359
 - l2tp unlock-test destination.....359
 - l2tp weighted-load-balancing command.....361
 - max-sessions command.....366, 371
 - sessions-limit-group command.....372
 - See also* show l2tp commands
- L2TP commands
 - l2tp
 - disconnect-on-local-pool-exhaustion.....376
 - l2tp reserve-vr.....373
 - session-virtual-router.....373
- L2TP dial-out
 - before configuring.....416
 - configuring.....417
 - dial-out process.....410
 - network model.....408
 - operational states.....411
 - outgoing call setup details.....411
 - Access-Accept message.....411
 - Access-Request message.....411
 - mutual authentication.....411
 - outgoing call successful.....411
 - route installation.....411

overview.....	407	applying through RADIUS.....	394
references.....	409	AVP relay, configuring.....	389, 391
route.....	408	configuration guidelines.....	389
session.....	408	configuring.....	391
target.....	408	how to apply.....	389
trigger.....	408	monitoring.....	444
L2tp dial-out commands.....	416	LAC (L2TP access concentrator).....	322
l2tp dial-out connecting-timer-value.....	417	before configuring.....	332, 366
l2tp dial-out dormant-timer-value.....	417	configuring receive window size (RWS).....	382
l2tp dial-out max-buffered-triggers.....	417	function.....	321, 331
l2tp dial-out session delete.....	417	sequence of events.....	323
l2tp dial-out session reset.....	417	Layer 2 Tunneling Protocol. <i>See</i> L2TP	
l2tp dial-out target.....	417	license commands.....	67
<i>See also</i> show l2tp dial-out commands		license b-ras.....	67
L2TP network server. <i>See</i> LNS		license l2tp-session.....	328
L2TP RWS (receive window size)		license service-manager.....	683
configuring global default.....	382	<i>See also</i> show license commands	
configuring on LAC.....	382	licenses	
configuring on LNS.....	382	B-RAS.....	67
l2tp tunnel default-receive-window		L2TP.....	328
command.....	382	Service Manager.....	683
overview.....	382	lifetime	
receive-window command (for LAC).....	382	guideline	
receive-window command (for LNS).....	382	for preferred lifetime.....	101
show l2tp command.....	122, 429, 431, 435	preferred	
show l2tp destination profile command.....	438	configuring for Prefix Delegation.....	101
L2TP transmit connect speed		restriction	
and Transmit (TX) Speed AVP 24.....	394	in configuration for delegated	
calculation methods		prefixes.....	101
how to configure.....	394	specifying	
monitoring.....	427, 430	as infinite.....	101
L2TP transmit connect speed and Transmit (TX)		valid	
Speed AVP 24		configuring for Prefix Delegation.....	101
calculation methods		limitation	
actual.....	397	on number of IPv6 prefixes	
dynamic layer 2.....	396	delegated to clients.....	103
examples.....	397	LLID (logical line identifier)	
QoS.....	396	configuration steps.....	31
static layer 2.....	396	how it works.....	31
configuring		monitoring.....	124, 132
AAA default tunnel parameters.....	400	preauthentication considerations.....	31
AAA domain maps.....	399	RADIUS attributes in preauthentication	
AAA tunnel groups.....	399	request.....	31
RADIUS.....	401	troubleshooting.....	31
reporting considerations.....	398	using to track subscribers.....	31
L2TP tunnel switch profiles		LNS (L2TP network server).....	323, 367
applying default profile.....	393	before configuring.....	332, 366
applying through AAA domain maps.....	392	call disconnect notification.....	376
applying through AAA tunnel groups.....	393	configuring.....	366

- configuring receive window size (RWS).....382
 - configuring VR.....373
 - installing multiple service modules.....377
 - modules supported.....379
 - out-of-resource result codes.....375
 - overriding out-of-resource result codes.....375
 - sequence of events.....323
 - local address pool
 - alias names.....26
 - ranges.....25
 - local address server.....25
 - alias names.....26
 - configuring.....25
 - pool ranges.....26
 - shared local address pools.....26
 - SNMP thresholds.....27
 - local authentication commands
 - aaa authentication default.....74
 - aaa local database.....74
 - aaa local select database.....74
 - aaa local username.....74
 - ip-address.....74
 - ip-address-pool74
 - operational-virtual-router.....74
 - password.....74
 - secret.....74
 - username.....74
 - local host command.....369
 - local ip address command.....369
 - logging. *See* specific feature, product, or protocol
 - logical line identifier, AAA. *See* LLID
 - logout subscribers command.....491
- M**
- MAC (media access control) addresses
 - duplicate.....490, 491
 - inheriting validation state.....632
 - preventing IP spoofing.....632
 - macros
 - service definitions.....664
 - Service Manager statistics.....718
 - manuals
 - comments on.....xxxv
 - maximum number of IPv6 prefixes
 - assigned to clients
 - common prefix for multiple subscribers.....40
 - unique prefix per subscriber.....40
 - using both DHCPv6 local server and Neighbor Discovery.....40
 - using Prefix Delegation only.....40
 - topologies in which they are assigned
 - same prefix for multiple subscribers.....40
 - unique prefix per subscriber.....40
 - MBS (RADIUS attribute 26-17).....213
 - media access control addresses. *See* MAC addresses
 - medium ipv4 command.....351, 354
 - merging policies
 - naming conventions.....673
 - Message-Authenticator (RADIUS attribute 80).....20
 - MLPPP Bundle Name (RADIUS attribute 26-62).....214
 - monitoring. *See* specific feature, product, or protocol
 - mutex service.....664, 691
- N**
- name server addresses
 - configuring.....24, 123
 - naming conventions
 - merged policies.....673
 - NAS (network access server).....305, 306
 - NAS-Identifier (RADIUS attribute 32).....210
 - NAS-IP-Address (RADIUS attribute 4).....210
 - NAS-Port (RADIUS attribute 5).....210
 - NAS-Port-Id (RADIUS attribute 87).....211
 - NAS-Port-Type (RADIUS attribute 61).....211
 - Ndra
 - enabling
 - IPv6 local address pool feature.....105
 - for IPv6 clients
 - overview.....57
 - guidelines for configuring
 - IPv6 local address pools.....58
 - NDRA
 - using IPv6 local address pools
 - monitoring a single pool.....173
 - monitoring all configured pools.....172, 174
 - Ndra clients
 - assigning prefixes to
 - using local address pools.....105
 - Ndra local address pools *See* IPv6 local address pools
 - network access server. *See* NAS

network commands	
network.....	653
non-PPP clients	
pool section for Prefix Delegation.....	104
non-PPP equal access	
configuration example.....	507
requirements.....	472
none domain name.....	8
notice icons.....	xxxiii

O

operational states, L2TP.....	411
chassis.....	411
sessions.....	411
targets.....	411
virtual router.....	411
option	
82.....	497, 498, 499, 512, 514, 515, 529, 530, 531, 533
out-of-resource result codes, LNS.....	375
Output-Gigawords (RADIUS attribute 53).....	211
override-user command.....	17

P

packet detection dynamic subscriber	
interfaces.....	631
packet fragmentation.....	323
packet mirroring.....	232
packets	
demultiplexing.....	619
transmitting.....	321
Partition-Accounting-Off messages.....	193
Partition-Accounting-On messages.....	193
password command.....	351, 354
PCP-Server-Name(RADIUS attribute 26-165).....	215
PCR (RADIUS attribute 26-15).....	213
peer.....	323
peer resynchronization.....	385
performance impact	
IPv4 and IPv6 services	
in a dual stack, combined.....	697
in a dual stack, independent.....	697
persistent tunnels, creating.....	379
PIB (Policy Information Base).....	47
platform considerations	
PPP.....	309
Point-to-Point Protocol. <i>See</i> PPP	
Policy Information Base. <i>See</i> PIB	
policy management on subscriber interfaces.....	625

PPP (Point-to-Point Protocol)	
accounting statistics for tunneled	
sessions.....	403
B-RAS service support.....	6
platform considerations.....	309
ppp commands	
ppp aaa-accounting-broadcast.....	89
ppp aaa-profile.....	31
PPP profile	
interim, broadcast, and policy-based	
accounting.....	89
PPP subscribers	
DHCPv6 local server: standalone mode	
delegation of prefixes.....	480
PPPoE remote circuit ID.....	520
Pppoe-Description (RADIUS attribute 26-24).....	213
PPPv6 subscribers	
assigning prefixes to	
using IPv6 local address pools.....	105
preauthentication	
AAA LLID.....	31
B-RAS users.....	8
preference.....	361
preference command.....	351, 354
preference order	
in allocation of prefixes	
to IPv6 clients.....	55
in assignment of DNS addresses	
to IPv6 clients.....	55
in determining local address pool	
for allocation of IPv6 prefixes.....	54
preferred lifetime	
for delegated prefixes	
configuring.....	101
default.....	101
setting	
without expiration.....	101
Prefix Delegation <i>See</i> DHCPv6 Prefix Delegation	
prefixes	
allocated to clients from	
interface configuration.....	54
IPv6 local address pools.....	54
RADIUS Access-Accept message.....	54
assigned length of /128	
in IPv6 local address pools.....	53
assigning to	
DHCPv6 clients.....	101
ndra clients.....	105

- configuring ranges
 - for delegation to clients.....101
 - delegating by
 - DHCPv6 local server.....101
 - ndra local server.....105
 - delegating to clients
 - over non-PPP links.....52
 - over PPP links.....52
 - excluded from
 - delegation to clients.....102, 106
 - excluding
 - range and individual ones102, 106
 - limitation on
 - number assigned to clients.....103
 - order of preference
 - in allocation to clients.....55
 - preferred and valid lifetimes
 - configuring for delegated ones.....101
 - prerequisite
 - for configuring IPv6 local address pools
 - for ndra.....105
 - for Prefix Delegation.....101
 - primary authentication/accounting RADIUS
 - server.....69, 92
 - primary IP interface.....624
 - privilege authentication, TACACS+.....305
 - profile commands
 - profile.....673
- Q**
- QoS (quality of service)
 - calculation method for L2TP transmit connect
 - speed.....396
 - on subscriber interfaces.....625
 - QoS commands
 - qos-parameter.....676
 - qos-profile.....673
 - QoS-Profile-Name (RADIUS attribute 26-26).....213
- R**
- RADIUS
 - LAG subscribers information.....45
 - RADIUS (Remote Authentication Dial-In User Service)
 - AAA failure.....39
 - accounting methods.....17
 - attribute descriptions.....17, 181, 241
 - attributes supported.....241
 - authentication and accounting servers.....17
 - authentication methods.....17
 - Calling-Station-Id formats supported.....210
 - change of authority messages.....227
 - CLI AAA messages.....208
 - client to server interaction.....17
 - configuring servers.....17
 - description.....180
 - direct server access.....17
 - disconnect messages.....227
 - EAP authentication.....17
 - IETF attributes supported.....241
 - Juniper Networks VSAs supported.....248
 - L2TP transmit connect speed.....401
 - L2TP tunnel switch profiles, applying.....394
 - message types supported.....182
 - RADIUS dynamic-request server.....227
 - round-robin server access.....17
 - server access.....17
 - server request processing limit.....17
 - Service Manager attributes.....687
 - Service Manager tags.....689
 - services.....180
 - traffic shaping for PPP over ATM
 - interfaces.....34
 - VSAs (vendor-specific attributes)
 - for dynamic IP interfaces.....34
 - formats.....248
 - RADIUS attributes
 - preference order and
 - allocation of prefixes to IPv6 clients.....55
 - RADIUS authentication
 - scenarios in which show subscribers command
 - output varies
 - IPv6 virtual router name contained in
 - Access-Accept message.....482
 - RADIUS configured on default
 - router.....482
 - RADIUS not configured on default
 - router.....482
 - virtual router name returned in
 - Access-Accept message.....482
 - radius commands.....17
 - radius acct-session-id-format.....210, 244
 - radius algorithm.....17
 - radius calling-station-delimiter.....210, 243
 - radius calling-station-format.....210, 243, 338
 - radius connect-info-format
 - command.....366, 374

radius connect-info-format		radius include framed-ipv6-route.....	211
l2tp-connect-speed.....	211	radius include framed-route.....	210
radius dsl-port-type.....	211, 245	radius include icr-partition-id.....	215
radius ethernet-port-type.....	211, 245	radius include ingress-policy-name.....	213
radius ignore atm-mbs.....	213	radius include input-gigapkts	213
radius ignore atm-pcr.....	213	radius include input-gigawords.....	211
radius ignore atm-scr.....	213	radius include interface-description.....	214
radius ignore atm-service-category.....	213	radius include ipv6-egress-policy-name.....	214
radius ignore egress-policy-name.....	213	radius include ipv6-ingress-policy-name.....	214
radius ignore framed-ip-netmask.....	210	radius include ipv6-local-interface.....	213
radius ignore ingress-policy-name.....	213	radius include ipv6-nd-ra-prefix.....	214
radius ignore ipv6-egress-policy-name.....	214	radius include ipv6-primary-dns.....	213
radius ignore ipv6-ingress-policy-name.....	214	radius include ipv6-secondary-dns.....	213
radius ignore virtual-router.....	213	radius include ipv6-virtual-router.....	213
radius include		radius include l2c-downstream-data.....	214
ANCP (L2C)-related Juniper Networks		radius include l2c-upstream-data.....	214
VSAs.....	215	radius include	
radius include access-loop-parameters.....	214	l2tp-ppp-disconnect-cause.....	381
radius include acct-authentic.....	210	radius include mlppp-bundle-name.....	214
radius include acct-delay-time.....	210	radius include nas-identifier.....	210
radius include acct-link-count.....	211	radius include nas-port.....	210
radius include acct-multi-session-id.....	210	radius include nas-port-id.....	211
radius include acct-session-id.....	210	radius include nas-port-type.....	211
radius include acct-session-id		radius include output-gigapkts.....	213
access-request.....	244	radius include output-gigawords.....	211
radius include acct-terminate-cause.....	36, 210	radius include pcp-server-name.....	215
radius include acct-tunnel-connection.....	211	radius include pppoe-description.....	213
radius include ascend-num-in-multilink.....	212	radius include profile-service-description.....	213
radius include called-station-id.....	210	radius include qos-profile-name.....	213
radius include calling-station-id.....	210	radius include tunnel-assignment-id.....	211
radius include class.....	210	radius include tunnel-client-auth-id.....	211
radius include connect-info.....	211	radius include tunnel-client-endpoint.....	211
radius include delegated-ipv6-prefix.....	212	radius include tunnel-interface-id.....	213
radius include dhcp-gi-address.....	214	radius include tunnel-medium-type.....	211
radius include dhcp-mac-address.....	214	radius include tunnel-preference.....	211
radius include dhcp-options.....	214	radius include tunnel-server-attributes.....	212
radius include		radius include tunnel-server-auth-id.....	211
downstream-calculated-qos-rate.....	214	radius include tunnel-server-endpoint.....	211
radius include ds-lite-tunnel-name.....	212	radius include tunnel-type.....	211
radius include dsl-forum-attributes.....	217	radius include	
radius include egress-policy-name.....	213	upstream-calculated-qos-rate.....	214
radius include event-timestamp.....	211	radius nas-identifier.....	210
radius include framed-compression.....	210	radius nas-port-format.....	210, 242
radius include framed-interface-id.....	211	radius nas-port-format extended atm.....	210
radius include framed-ip-add acct-start.....	242	radius nas-port-format extended	
radius include framed-ip-addr.....	210	ethernet.....	210
radius include framed-ip-netmask.....	36, 210	radius override calling-station-id	
radius include framed-ipv6-pool.....	211	remote-circuit-id.....	210
radius include framed-ipv6-prefix.....	211	radius override nas-info.....	210

- radius override nas-ip-addr
 - tunnel-client-endpoint.....210
- radius override nas-port-id
 - remote-circuit-id.....211
- radius pppoe nas-port-format
 - unique.....210, 242
- radius pre-authentication server.....31
- radius relay server.....239
- radius relay udp-checksum.....239
- radius remote-circuit-id-delimiter.....210
- radius remote-circuit-id-format.....210
- radius route-download server.....29
- radius trap acct-server-not-responding.....71
- radius trap acct-server-responding.....71
- radius trap auth-server-not-responding.....71
- radius trap auth-server-responding.....71
- radius trap no-acct-server-responding.....71
- radius trap no-auth-server-responding.....71
- radius update-source-addr.....180, 242
- radius vlan nas-port-format stacked.....242
- See also show radius commands
- radius dynamic-request
 - platform.....228
- RADIUS dynamic-request server
 - change of authorization messages.....232
 - disconnect messages.....229
 - how it works229
 - message exchange.....229, 232
 - overview.....227
 - qualifications for disconnect.....229
 - security and authentication.....229
 - Service Manager.....731
- RADIUS IPv6 attributes
 - configuring
 - for DHCPv6 Prefix Delegation.....39
 - for IPv6 Neighbor Discovery.....39
 - verifying
 - for DHCPv6 Prefix Delegation.....154
 - for IPv6 Neighbor Discovery.....154
- radius relay
 - platform.....238
- RADIUS Relay Server.....235
 - SRC Software.....238
- RADIUS relay server
 - configuring.....239
- radius remote-circuit-id-format command.....512
- RADIUS route-download server.....29
 - configuring.....29
 - format of routes.....29
 - how it works.....29
 - per chassis.....29
 - supported attributes.....29
- RADIUS servers
 - assignment of a unique prefix route
 - to each IPv6 client.....41
 - total number of routes used for
 - delegation.....41
 - Prefix Delegation and
 - pool name not returned in
 - Access-Accept.....54
 - pool name returned in Access-Accept.....54
- radius-attributes-override-monitoring.....291
- RADIUS-initiated change of authorization
 - qualifications for change of authorization.....232
- RADIUS-initiated COA
 - configuring.....234
- RADIUS-initiated disconnect
 - configuring.....231
 - L2TP LAC users.....231
 - references.....229
 - sample network.....228
 - security and authentication.....232
- realm names
 - configuring.....9
 - usage.....9
- Receive speed AVP.....435
- receive window size (RWS). See L2TP RWS
- redirected authentication.....8
- remote access (B-RAS). See B-RAS applications
- Remote Authentication Dial-In User Service. See RADIUS
- remote clients, IP addresses for.....4
- remote host command.....369
- remote system.....323
- requesting routers
 - as customer edge device
 - in obtaining IPv6 prefixes.....51
 - assigning prefixes to
 - using IPv6 local address pools.....101
 - receipt of IPv6 prefixes
 - from delegating routers.....51
- router-name command.....351, 354
- RX Connect-Speed AVP.....354
- RX speed AVP.....354, 355

S**S-VLAN links**

- between CPE and PE routers
- pool section for Prefix Delegation.....104

SCR (RADIUS attribute 26-16).....213

SDX (Service Deployment System) software.....139

See also SRC software

Sending RADIUS Attributes

Overview.....495

server-name command.....351, 354

service commands

service dhcp-external550

service dhcp-local.....491, 653

service dhcp-local authenticate.....497

service dhcpv6-local501

service dhcpv6-local standalone
authenticate.....500

service definitions.....663, 664, 667

copying.....670

creating.....667

installing.....670

modifying.....670

modifying QoS configurations.....676

specifying parameter instances.....673

specifying QoS profiles.....673

uninstalling.....670

service dhcp-local command.....499

service dhcpv6-local command.....501

service instance.....664

service interim accounting

in a dual stack

of IPv4 and IPv6 services.....705

IPv4 and IPv6 services

overview.....705

Service Manager

CLI support.....684

COA-Request messages.....731

combined IP4 and IPv6 service

example.....740

configuring

Service Manager license.....683

deactivating.....689

setting thresholds.....689

dual-stack subscribers

IPv6 service accounting statistics, separate
collection.....706

guided entrance.....663, 730, 731

IPv4 and IPv6 services

combined, activating.....696

combined, deactivating.....696

combined, overview.....694

in a dual stack, activating.....696

in a dual stack, overview.....694

independent, activating.....696

independent, deactivating.....696

independent, overview.....694

overview.....694, 696

license sessions.....683

macros.....664

multiple services.....692

mutex service.....664, 691

overview.....663

parameter values.....688

preprovisioning services.....707, 711

QoS

considerations.....681

modifying configurations of676

referencing configurations of.....673

removing references of.....676

RADIUS dynamic-request server.....731

RADIUS support.....684

RADIUS tags.....689

service definition.....663, 664, 667

parameters.....708

service instance.....664

service session.....664

forcing deactivation.....715

profiles.....711

service session profiles.....664

statistics.....688, 711, 718

macro command.....718

using RADIUS.....720

using the CLI.....721

statistics collection

for external parent groups, setting

up.....722

subscriber session ID.....717

subscriber sessions.....707

supported platforms.....664

tasks.....665

testing services.....707

Service Manager commands

no service-management owner-session

force.....715

no service-management subscriber-session

force.....715

- service-management install.....670
- service-management owner-session.....707
- service-management
 - service-session-profile.....711
- service-management subscriber-session
 - service-session.....707
- statistics.....711
- time.....711
- volume.....711
- Service Manager license
 - configuring.....683
- service modules
 - installing multiple for LNS sessions.....377
- service session.....664
- Service-Category (RADIUS attribute 26-14).....213
- Service-Description (RADIUS attribute 26-53).....213
- Service-Interim-Acct-Interval (RADIUS attribute 26-140).....687, 698
- Service-Session (RADIUS attribute 26-83).....687, 698
- Service-Statistics (RADIUS attribute 26-69).....687
- Service-Stats (RADIUS attribute 26-69).....720
- Service-Timeout (RADIUS attribute 26-68).....687, 690
- Service-Volume (RADIUS attribute 26-67).....687, 690
- session.....323
- Session and Resource Control. *See* SRC software
- session timeout
 - configuring.....38
 - interpreting default value.....38
 - range for.....38
- session-out-of-resource-result-code-override
 - command.....375
- Session-Timeout (RADIUS attribute 27).....20
- sessions, L2TP.....328
- set dhcp commands
 - set dhcp vendor-option.....512
- set dhcp relay commands
 - set dhcp relay.....512
 - set dhcp relay agent sub-option512
 - set dhcp relay assign-giaddr-source-ip.....512
 - set dhcp relay broadcast-flag-replies.....512, 516
 - set dhcp relay giaddr-selects-interface.....653
 - set dhcp relay layer2-unicast-replies.....512, 521
 - set dhcp relay options.....512
 - set dhcp relay override.....512
 - set dhcp relay
 - preserve-trusted-client-option.....512
 - set dhcp relay proxy.....543
 - set dhcp relay trust-all.....512
- set dhcp relay proxy commands
 - set dhcp relay proxy.....543
 - set dhcp relay proxy send-first-offer.....543
 - set dhcp relay proxy timeout.....543
- set dhcp vendor-option command.....523
- shared IP interfaces.....624
- shared local address pools.....26
- shared tunnel-server ports.....326, 377
 - using with L2TP.....332, 367
- show aaa commands
 - show aaa accounting.....115
 - show aaa accounting default.....118
 - show aaa accounting interval.....118
 - show aaa accounting vr-group.....118
 - show aaa authentication default.....119
 - show aaa dhcpv6-delegated-prefix.....154
 - show aaa domain-map.....120, 122, 427
 - show aaa duplicate-address-check.....123
 - show aaa intf-desc-format.....295
 - show aaa ipv6-nd-ra-prefix.....154
 - show aaa model.....123
 - show aaa name-servers.....123
 - show aaa profile.....125
 - show aaa
 - radius-override-ncp-negotiation.....175
 - show aaa route-download ipv6.....125
 - show aaa route-download ipv6 routes.....127
 - show aaa route-download ipv6 routes
 - global.....129
 - show aaa route-download routes.....127
 - show aaa route-download routes global.....129
 - show aaa send-delay-start.....117
 - show aaa send-immediate-update-ipv6.....117
 - show aaa service accounting interval.....753
 - show aaa statistics.....131
 - show aaa subscriber per-port-limit.....133
 - show aaa subscriber per-vr-limit.....133
 - show aaa timeout.....134
 - show aaa tunnel-group.....427, 430
 - show aaa tunnel-parameters.....430, 432
 - show aaa user accounting interval.....134
 - show radius route-download.....125
- show configuration commands
 - show configuration category aaa
 - global-attributes.....134
 - show configuration category aaa
 - local-authentication.....135

show configuration category aaa		
server-attributes include-defaults.....	136	
show configuration category		
aaaglobal-attributes.....	134	
show configuration category		
aalocal-authentication.....	135	
show configuration category		
aaaserver-attributes include-defaults.....	136	
show cops info command.....	138	
show cops statistics command.....	140	
show dhcp commands		
show dhcp binding.....	562	
show dhcp count.....	565	
show dhcp host.....	567	
show dhcp proxy-client binding.....	600	
show dhcp relay.....	585	
show dhcp relay proxy statistics.....	586	
show dhcp relay statistics.....	588	
show dhcp server.....	592	
show dhcp server statistics.....	591	
show dhcp summary command.....	599	
show dhcp vendor-option command.....	583	
show dhcp-external command.....	574	
show ip commands		
show ip demux interface.....	655	
show ip local alias.....	142	
show ip local pool.....	143	
show ip local-pool statistics command.....	144	
show ip service-profile.....	615	
show ip-subscriber.....	616, 656	
show ip dhcp-capture command.....	584	
show ip dhcp-external commands		
show ip dhcp-external binding.....	570	
show ip dhcp-external client-id.....	570	
show ip dhcp-external configuration.....	572	
show ip dhcp-external statistics.....	573	
show ip dhcp-local commands		
show ip dhcp-local.....	578	
show ip dhcp-local auth.....	577	
show ip dhcp-local binding.....	571	
show ip dhcp-local excluded.....	561	
show ip dhcp-local limits.....	597	
show ip dhcp-local pool.....	575	
show ip dhcp-local reserved.....	599	
show ip dhcp-local statistics.....	579, 580	
show ip http commands		
show ip http scalar.....	750	
show ip http server.....	750	
show ip http statistics.....	751	
show ip interface commands		
show ip interface.....	755	
show ip local shared-local command.....	145	
show ipv6 dhcpv6-local commands		
show ipv6 dhcpv6-local auth.....	596	
show ipv6 dhcpv6-local binding.....	593	
show ipv6 dhcpv6-local		
dns-domain-searchlist.....	594	
show ipv6 dhcpv6-local dns-servers.....	594	
show ipv6 dhcpv6-local limits.....	600	
show ipv6 dhcpv6-local prefix-lifetime.....	595	
show ipv6 dhcpv6-local statistics.....	595	
show ipv6 dhcpv6-local ssc		
show aaa dhcpv6-delegated-prefix.....	56, 602	
show ipv6 interface commands		
show ipv6 interface.....	755	
show ipv6 local pool commands		
for a single pool.....	170	
for all configured pools.....	169, 172, 173, 174	
statistics for a single pool.....	171	
show l2tp commands		
show l2tp.....	433	
show l2tp destination.....	435	
show l2tp destination lockout.....	437	
show l2tp destination summary.....	440	
show l2tp session.....	442	
show l2tp session summary.....	443	
show l2tp tunnel.....	445	
show l2tp tunnel summary.....	448	
show l2tp dial-out commands		
show l2tp dial-out.....	449	
show l2tp dial-out target.....	455	
show license commands		
show license b-ras.....	146	
show license service-management.....	753	
show profile commands		
show profile	754	
show profile name.....	752	
show radius commands		
show radius accounting servers.....	147	
show radius accounting statistics.....	149	
show radius algorithm.....	146	
show radius attributes-included.....	298	
show radius authentication servers.....	147	
show radius authentication statistics.....	149	
show radius calling-station-delimiter.....	293	
show radius calling-station-format.....	293	
show radius connect-info-format.....	295	
show radius dsl-port-type.....	294	

- show radius dynamic-request servers.....300
- show radius dynamic-request statistics.....299
- show radius ethernet-port-type.....294
- show radius icr-partition-accounting.....304
- show radius nas-identifier.....293
- show radius nas-port-format.....292
- show radius override.....291
- show radius pppoe nas-port-format.....292
- show radius remote-circuit-id-delimiter.....294
- show radius remote-circuit-id-format.....294
- show radius rollover-on-reject.....147
- show radius route-download statistics.....147
- show radius servers.....147, 300
- show radius statistics.....147, 149, 299
- show radius trap.....153
- show radius tunnel-accounting.....153
- show radius update-source-address.....154
- show radius vlan nas-port-format.....292
- show radius override.....291
- show radius relay commands
 - show radius relay servers.....303
 - show radius relay statistics.....301
 - show radius relay udp-checksum.....303
- show service-management commands
 - show service-management
 - owner-session.....767
 - show service-management
 - service-definition.....765
 - show service-management
 - service-session-profile.....766
 - show service-management
 - subscriber-session.....770
 - show service-management summary.....773
- show sccc commands
 - show sccc info.....155
 - show sccc option.....160
 - show sccc statistics.....158
 - show sccc version.....159
- show subscribers command.....161
- show terminate-code command.....167
- SNMP traps
 - configuring for DHCP local servers.....490
 - configuring for RADIUS servers.....71
- source-address command.....351, 354
- spoofing, IP
 - preventing.....632
- SRC (Session and Resource Control)
 - software.....138, 140, 461, 507
 - configuring the client.....47
 - monitoring the client.....155, 158, 160
- SRC (Session and Resource Control)software.....138
- sscc commands.....47
 - sscc address.....47
 - sscc enable.....47
 - sscc protocol ipv6.....47
 - sscc protocol lac.....47
 - sscc retryTimer.....47
 - sscc sourceAddress.....47
 - sscc transportRouter.....47
 - See also show sccc commands
- standalone DHCP local server.....474
- standalone mode
 - with and without AAA authentication
 - DHCPv6 local server functionality.....480
- standalone mode with AAA authentication
 - DHCPv6 local server
 - accounting of IPv6 subscribers.....481
 - DHCPv6 local server, parameters in username
 - circuit identifier.....480
 - circuit type.....480
 - domain name.....480
 - user prefix.....480
 - virtual router.....480
- standard RADIUS attributes
 - configuring
 - for DHCPv6 Prefix Delegation.....39
 - for IPv6 Neighbor Discovery.....39
 - IPv6 Neighbor Discovery and
 - configuring logging severity.....39
 - warning message.....39
 - using the same values
 - for Neighbor Discovery and Prefix Delegation.....39
 - verifying
 - for DHCPv6 Prefix Delegation.....154
 - for IPv6 Neighbor Discovery.....154
- State (RADIUS attribute 24).....20
- statistics
 - for DHCPv6 Prefix Delegation
 - viewing.....171
 - for NDRA
 - viewing.....174
- statistics collection
 - for external parent groups
 - setting up.....722

subscriber interface commands	
set dhcp relay giaddr-selects-interface.....	653
subscriber interfaces	
applications.....	626
commands	
configuring dynamic.....	653
configuring static.....	635
configuring.....	628
multicast routing protocols.....	625
policies and QoS.....	625
routing protocols.....	625
dynamic.....	642
inheriting MAC validation state.....	632
GRE tunnel configuration.....	648
IP over bridged Ethernet configuration.....	647
IP over Ethernet configuration.....	642
IP over VLAN over Ethernet	
configuration.....	645, 649
monitoring.....	655
overview	
dynamic.....	628
static.....	624
static.....	635
subscribers	
accounting messages.....	192
authorization and authentication	
messages.....	182
E Series Broadband Services Routers.....	161
limiting active subscribers.....	39
preauthentication and AAA LLID.....	31
support, technical See technical support	
T	
TACACS+	
AAA services.....	305
accounting.....	305
authentication login process.....	305
authorization.....	305
configuring.....	310
daemon.....	305, 306
host.....	306
NAS (network access server).....	305, 306
privilege authentication.....	305
TACACS+ commands	
aaa accounting commands.....	310
aaa accounting exec.....	310
aaa accounting suppress null-username.....	310
aaa authentication enable default.....	305, 310
aaa new-model.....	305
TCP and TACACS+.....	305
technical support	
contacting JTAC.....	xxxvi
Terminal Access Controller Access Control System	
+. See TACACS+	
text and syntax conventions.....	xxxiv
timeout, configuring for B-RAS applications.....	38
traffic shaping for PPP over ATM.....	36
transmit connect speed, L2TP. See L2TP transmit	
connect speed	
tunnel	
defined.....	321, 323
selection, L2TP.....	361
switching.....	379
tunnel commands, L2TP	
tunnel.....	348, 352
tunnel group.....	348
tunnel password.....	366
tunnel group mode, mapping to L2TP tunnel.....	352
tunnel selection, L2TP.....	361
failover between preference levels.....	361
failover within preference levels.....	361
maximum sessions per tunnel.....	361
weighted load balancing.....	361
tunnel subscribers, enabling authentication for.....	23
tunnel switch profiles, L2TP	
applying default profile.....	393
applying through AAA domain maps.....	392
applying through AAA tunnel groups.....	393
applying through RADIUS.....	394
AVP relay, configuring.....	389, 391
configuration guidelines.....	389
configuring.....	391
how to apply.....	389
monitoring.....	444
Tunnel-Assignment-Id (RADIUS attribute 82).....	211
Tunnel-Client-Auth-Id (RADIUS attribute 90).....	211
Tunnel-Client-Endpoint (RADIUS attribute	
66).....	211
Tunnel-Interface-Id (RADIUS attribute 26-44).....	213
Tunnel-Medium-Type (RADIUS attribute 65).....	211
Tunnel-Preference (RADIUS attribute 83).....	211
tunnel-server command.....	367
tunnel-server ports	
shared.....	326
Tunnel-Server-Auth-Id (RADIUS attribute 91).....	211
Tunnel-Server-Endpoint (RADIUS attribute	
67).....	211
tunnel-subscriber authentication command.....	24

Tunnel-Type (RADIUS attribute 64).....	211
tunneled PPP session accounting statistics.....	403
tunnels, IP	
shared tunnel-server ports.....	326
tx-connect-speed-method command.....	402
Type 1 DUID	
not supported by DHCPv6 local server.....	504
Type 2 DUID	
default type present on DHCPv6 local	
server.....	504
Type 3 DUID	
configured using ipv6 dhcpv6-local duid-type	
command.....	504
type command, L2TP.....	348, 352

U

UDP (User Datagram Protocol)	
checksums.....	22, 153
Upstream-Calculated-QoS-Rate (RADIUS attribute	
26-142).....	214
user domain, mapping to L2TP tunnel.....	348
User-Name (RADIUS attribute 1).....	9
usernames and passwords from a domain	
configuring.....	16
using shared tunnel-server ports.....	367

V

valid lifetime	
for delegated prefixes	
configuring.....	101
default.....	101
setting	
without expiration.....	101
validation state, inheriting for MAC addresses	
.....	632
vendor class identifier option.....	522
vendor-specific (suboption 9).....	526
vendor-specific attributes. <i>See</i> VSAs	
virtual router group	
interim, broadcast, and policy-based	
accounting.....	89
virtual routers	
mapping user domain names.....	7, 120, 131, 427
redirected authentication.....	8
Virtual-Router (RADIUS attribute 26-1).....	213
vlan commands	
vlan advisory-rx-speed.....	356
vlan id.....	653

VLAN links	
between CPE and PE routers	
pool section for Prefix Delegation.....	104
VLANs (virtual local area networks)	
configuring dynamic subscriber	
interfaces.....	645, 649
VPNs (virtual private networks)	
connecting subscribers	627
VSAs (vendor-specific attributes)	
DSL Forum	
controlling inclusion of.....	217
descriptions.....	260
in AAA access and accounting	
messages.....	206
for dynamic IP interfaces.....	34
formats.....	248

W

walled garden. <i>See</i> guided entrance	
Web access to E Series router.....	461
Windows Internet Name Service. <i>See</i> WINS	
WINS, assigning IP addresses.....	24, 123

