

Stinger[®] IP Control Module

Configuration Guide

Copyright © 2007 Alcatel-Lucent. All Rights Reserved.

This material is protected by the copyright and trade secret laws of the United States and other countries. It may not be reproduced, distributed, or altered in any fashion by any entity (either internal or external to Alcatel-Lucent), except in accordance with applicable agreements, contracts, or licensing, without the express written consent of Alcatel-Lucent and the business management owner of the material.

Trademarks

All trademarks and service marks specified herein are owned by their respective companies.

Customer Service

Product and service information, and software upgrades, are available 24 hours a day. Technical assistance options accommodate varying levels of urgency.

Finding information and software

To obtain software upgrades, release notes, and addenda for this product, log in to Alcatel-Lucent OnLine Customer Support at <http://www.lucent.com/support>.

Alcatel-Lucent OnLine Customer Support also provides technical information, product information, and descriptions of available services. The center is open 24 hours a day, seven days a week. Log in and select a service.

Obtaining technical assistance

Alcatel-Lucent OnLine Customer Support at <http://www.lucent.com/support> provides easy access to technical support. You can obtain technical assistance through email or the Internet, or by telephone. If you need assistance, make sure that you have the following information available:

- Active service or maintenance contract number, entitlement ID, or site ID
- Product name, model, and serial number
- Software version or release number
- Software and hardware options
- If supplied by your carrier, service profile identifiers (SPIDs) associated with your line
- Whether you are routing or bridging with your Alcatel-Lucent product
- Type of computer you are using
- Description of the problem

Obtaining assistance through email or the Internet

If your services agreement allows, you can communicate directly with a technical engineer through Email Technical Support or a Live Chat. Select one of these sites when you log in to <http://www.lucent.com/support>.

Calling the technical assistance center (TAC)


The Alcatel-Lucent Technical Assistance Center (TAC) is available to assist you with any problems encountered while using this Alcatel-Lucent product. Call the appropriate number as given below:

Inside USA: +1-866-582-3688

Outside USA: +1-630- 218-7688

You can also log on to our Customer Support web site:
<http://www.alcatel-lucent.com/support>

Contents



About This Guide	xix
Chapter 1 Introduction	1-1
Stinger IP DSLAM network features	1-1
Introduction to the Stinger IP2000 control module	1-3
IP2000 model numbers and platform support	1-3
IP2000 support for up to 2048 trunk terminated calls	1-4
Features not currently supported by the IP2000	1-5
Introduction to the Stinger IP2100 control module	1-5
Model numbers and platform support	1-6
IP2100 memory modes	1-7
Maximizing the number of user connections	1-7
Overview of OAM-CONFIG profile settings	1-8
Required IP2100 memory mode to maximize user connections	1-8
Additional connections created for trunk-terminated connections	1-8
Notice of OAM limitation with more than 5000 user PVCs	1-8
Configurable IP2100 port managers	1-9
System-generated profiles	1-9
Fixed-rate and nonfixed-rate bandwidth	1-9
Configurable NP-PORT profile settings	1-10
Caution about modifying NP-PORT bandwidth allocations	1-12
How to recover from a bandwidth starvation condition	1-13
Protection against exceeding line capacity in NP-PORT configurations	1-13
Sample HB LIM NP-PORT configuration	1-13
Features not currently supported by the IP2100	1-15
Network architecture overview	1-15
Multicast video	1-15
Internet and voice access	1-16
Multiplexing multiple IP flows on a single ATM VCC	1-17
Chapter 2 Gigabit Ethernet Configuration	2-1
Configuring Gigabit Ethernet interfaces	2-1
Overview of ETHERNET profile settings	2-1
Enabling layer 2 bridging for VLAN operations	2-3
Verifying the Gigabit Ethernet interface setup	2-3
Checking the routing table	2-4
Verifying the network processor setup for the interface	2-4
Verifying the SAR setup for the interface	2-4
Verifying IP packet transfer on the interface	2-4
Configuring Gigabit Ethernet port redundancy	2-5

Configuring a soft IP interface for Gigabit Ethernet redundancy	2-6
Configuring Gigabit Ethernet redundancy for RFC 2684 (IPoA) connections ..	2-6
Configuring Gigabit Ethernet redundancy for VLAN bridging	2-7
Configuring a redundant LAN MBONE.....	2-8
Configuring LACP on Gigabit Ethernet ports (IP2100 only).....	2-9
LACP configuration overview.....	2-10
ETHER-GROUP settings	2-10
Virtual ETHERNET profile settings for a LAG	2-10
ETHERNET profile settings for physical ports in a LAG	2-12
LACP profile settings	2-13
Sample link aggregation configuration	2-14
Configuring a routed VLAN on the aggregated bandwidth.....	2-14
Configuring an N:1 bridged VLAN that uses the aggregated bandwidth ..	2-15
Configuring an MBONE interface on the aggregated ports.....	2-15
LACP implementation details	2-16
LACP-related diagnostics	2-17
Configuring STP on Gigabit Ethernet ports (IP2100 only).....	2-17
Primary application with this software version	2-17
Limitations with this software version	2-18
Configuration overview	2-18
Bridge-level STP settings in the ETHER-GROUP profile	2-19
Port-level STP settings in each port's ETHERNET profile.....	2-20
Sample configuration with transparent bridging.....	2-22
Routing implications for STP-enabled ports	2-24
Administrative tools for Gigabit Ethernet	2-25

Chapter 3 VLAN Configuration.....3-1

IP filters for Ethernet-encapsulated bridged IP datagrams	3-2
Configuring 1:1 VLAN bridging.....	3-2
Overview of VLAN-ETHERNET and CONNECTION settings.....	3-3
Sample 1:1 VLAN bridging configuration	3-5
Configuring N:1 VLAN bridging	3-6
Creating and configuring bridge groups	3-6
Overview of BRIDGE-GROUP settings.....	3-7
Sample BRIDGE-GROUP configuration with MAC address aging	3-9
Sample BRIDGE-GROUP configuration with port blocking.....	3-9
Sample BRIDGE-GROUP configuration with IGMP snooping.....	3-10
VLAN and connection settings.....	3-11
How address limiting works.....	3-11
Sample N:1 VLAN bridging configuration with address limiting.....	3-12
Configuring DHCP and PPPoE snooping for DSL line identification	3-14
Structure of identifiers added to DHCP and PPPoE Discovery packets	3-14
Configuring DHCP snooping	3-16
Configuring PPPoE snooping for line identification of PPPoE clients	3-20
Dynamic anti-spoofing for DHCP and PPPoE	3-24
IP anti-spoofing through DHCP.....	3-25
PPPoE anti-spoofing	3-25
Configuring dynamic anti-spoofing.....	3-25
Anti-spoofing aging.....	3-25
Configuring stacked VLANs.....	3-27
Bridging untagged frames to stacked VLANs.....	3-28

Overview of VLAN stacking settings for untagged frames	3-28
Sample configuration bridging untagged frames.....	3-29
Bridging enterprise VLAN tagged frames to stacked VLANs.....	3-30
Overview of VLAN stacking settings for tagged frames.....	3-31
Sample configuration for mapping tagged frames.....	3-33
Configuring routed VLANs.....	3-35
Creating a virtual IP interface for a routed VLAN	3-35
Sample routed VLAN configuration	3-36
Applying an IP filter to a routed VLAN	3-38
Assigning a virtual router to a routed VLAN.....	3-38
VLAN bridging of IPoA traffic to an upstream BRAS	3-39
Upstream packet processing.....	3-39
Downstream packet processing	3-40
Configuration overview	3-40
Overview of the BRAS profile.....	3-40
Overview of CONNECTION profile settings	3-41
Configuring VLAN bridged IPoA.....	3-42
Configuring bridged IPoA to multiple BRAS with the same IP address	3-44
Caveat when configuring an unnumbered interface for bridged IPoA.....	3-46
Sample VLAN configuration.....	3-46
Sample profiles	3-46
Possible proxy ARP problems with the sample configuration	3-47
Conclusions about the sample configuration	3-48
VLAN-based ARP table	3-48
How the VLAN ARP entries are updated dynamically	3-48
Adding VLAN ARP entries statically.....	3-48
PPPoA bridging in 1:1 and N:1 VLAN	3-49
Upstream network element configuration requirements	3-49
PPPoA bridging in N:1 VLANs	3-50
Required settings to enable PPPoA bridging in a bridge group	3-50
Downstream CPE configuration considerations.....	3-51
Sample configuration enabling PPPoA bridging in a bridge group	3-51
PPPoA bridging in 1:1 VLANs.....	3-53
Required setting to enable PPPoA bridging in a 1:1 VLAN	3-53
Downstream CPE configuration considerations.....	3-54
Sample configuration enabling PPPoA bridging in a 1:1 VLAN circuit	3-54
Administrative tools for VLAN.....	3-56
 Chapter 4 IP Routing Configuration.....	 4-1
Introduction to the IP router software	4-1
Routes and interfaces.....	4-2
Displaying the routing table	4-2
Displaying the interface table	4-3
IP control module performance statistics	4-4
IP address syntax	4-4
Configuring IP-INTERFACE profiles for Ethernet ports.....	4-6
Overview of typical local interface settings	4-6
Configuring a local IP interface.....	4-7
Defining a local virtual IP interface	4-8
Defining a soft interface for increased accessibility.....	4-8
Disabling directed broadcasts to protect against denial-of-service	4-8

Configuring IP-GLOBAL network features	4-9
Setting a system address	4-9
Configuring DNS.....	4-10
Overview of typical DNS settings	4-10
Specifying domain names for lookups	4-10
Setting RIP options	4-11
RIP policy for propagating updates back to the originating subnet.....	4-12
RIP triggering	4-12
Limiting the size of UDP packet queues.....	4-12
Ignoring default routes when updating the routing table.....	4-13
Suppressing host-route advertisements	4-13
Subscriber profile sharing	4-14
Equal cost multipath routing.....	4-15
Packet-based ECMP	4-15
Flow-based ECMP	4-15
Caveat about changing the type of load balancing	4-15
Limitations of ECMP load balancing	4-16
Overview of ECMP settings.....	4-16
Sample downstream ECMP configuration	4-17
Example of downstream flow-based ECMP	4-19
Sample upstream ECMP through routed VLAN interfaces	4-21
Configuring and using address pools.....	4-23
Pool requirements	4-23
Overview of settings for defining pools.....	4-23
Preventing the use of class boundary addresses.....	4-26
Examples of configuring address pools	4-26
Example of configuring summarized address pools	4-27
Examples of assigning an address from a pool	4-29
IP pool chaining	4-30
Slot-based address assignment.....	4-36
Overview of profile settings.....	4-36
Sample slot-based address configuration using the global router	4-37
Expanding the sample configuration to use a virtual router	4-39
Configuring IP-ROUTE profiles	4-40
Overview of typical static route settings	4-40
Offloading routing overhead to an external router	4-41
Creating a static route to a subnet.....	4-42
Overview of routed subscriber connection features.....	4-42
Source interface local addresses.....	4-42
Packets that use the specified source address.....	4-42
CPE client considerations	4-43
Soft IP interface requirement	4-43
Overview of configuration settings	4-44
Sample configuration with a source interface address.....	4-44
Anti-spoofing protection for IPoA, BIR, PPPoA, and PPPoE connections	4-47
Overview of anti-spoofing settings.....	4-47
Sample anti-spoofing configuration	4-48
Configuring IPoA subscriber connections	4-49
Typical CONNECTION atm-options settings for terminating PVCs	4-49
Typical CONNECTION ip-options settings for terminating PVCs	4-50
Sample RFC 2684 (IPoA) terminating PVC	4-51
Example of using a local-address setting for a numbered interface	4-52

Example of routing a terminated PVC across Gigabit Ethernet	4-53
Example of using IP routing to aggregate PVCs onto a trunk VC.....	4-55
Configuring BIR subscriber connections	4-56
Overview of CONNECTION bir-options and ip-options settings	4-56
Sample subnet (BIR/24) configuration.....	4-57
Sample host route (BIR/32) configurations.....	4-58
Sample BIR connection on a VDSL port.....	4-60
Sample use of filters with BIR connections	4-61
Configuring multiple WAN virtual IP interfaces on a BIR connection	4-62
Overview of multiple IP address settings	4-63
Sample configuration of multiple addresses on a BIR connections	4-64
Configuring DHCP relay for IPoA and BIR connections	4-65
RFC compliance and caveats	4-65
DHCP option 82.....	4-65
WAN virtual interfaces on DHCP-configured BIR connections.....	4-66
IP address assignments on DHCP-configured WAN virtual interfaces	4-66
DHCP IP address lease time for WAN virtual interfaces.....	4-66
When WAN virtual interfaces are activated.....	4-67
How the system selects an interface for incoming packets	4-67
Deactivating and deleting WAN virtual IP interfaces.....	4-67
DHCP relay configuration settings.....	4-68
Overview of IP-GLOBAL DHCP relay settings	4-68
Overview of DHCP relay agent settings.....	4-68
Overview of DHCP option 82 suboption settings.....	4-70
DHCP option 82 config sub-options 1 and 2	4-71
Overview of IP-INTERFACE and CONNECTION DHCP settings.....	4-72
Per-connection control of the DHCP relay giaddr field	4-73
Overview of how to use per-connection giaddr contents	4-73
Samples of how various settings affect giaddr contents	4-74
Sample DHCP relay configurations for IPoA connections	4-74
Sample configuration using DHCP relay without option 82	4-74
Sample configuration using DHCP relay with option 82	4-75
Sending only a hostname in the suboption fields	4-75
Interoperation with DHCP servers that zero-delimit suboption fields	4-76
Allowing non-standard DHCP source ports	4-76
DHCP issues on LAN management interfaces	4-76
Sample DHCP relay configurations for BIR connections.....	4-78
Sample configuration with no DHCP relay on a BIR connection	4-78
Sample configuration enabling relay agent on a BIR connection	4-79
Sample configuration with option 82	4-80
Sample configuration with option 82 and multiple interface creation	4-81
Sample configuration on a VDSL port with option 82 and multiple interface creation.....	4-82
Sample using the soft interface address in the giaddr field	4-82
Sample configuration using the DHCP router option	4-83
Sample configuration that sends only a hostname	4-84
Configuring broadband RAS subscriber access.....	4-85
Recommended call-type setting for PPP sessions	4-85
Overview of PPPoA and PPPoE topologies	4-86
Required setup for PPPoA and PPPoE connections	4-87
Configuring the ANSWER-DEFAULTS profile for PPP sessions	4-87
Terminating traffic on a LIM internal interface	4-89

The underlying ATM circuit required to terminate traffic	4-89
PPPoA and PPPoE encapsulation types in the underlying ATM circuit	4-89
Configuring an ATM circuit to terminate PPPoA	4-90
Configuring an ATM circuit to terminate PPPoE	4-91
Example of configuring a PPPoA connection	4-91
Overview of PPPoA CONNECTION settings.....	4-92
Sample PPPoA connection with bidirectional CHAP authentication	4-93
PPPoA over LLC.....	4-94
Verifying the encapsulation type.....	4-94
Example of configuring a PPPoE connection.....	4-95
Overview of PPPoE CONNECTION settings	4-95
Sample PPPoE connection using PAP authentication	4-96
Sample PPPoE connection on a VDSL port	4-97
Enabling LCP keepalives	4-98
Enabling magic-number negotiation when LQM is disabled	4-101
Enabling support for PPPoE service names.....	4-101
Details of the service-name negotiation process	4-101
Configuration overview	4-102
Sample PPPoE server configuration using local profiles.....	4-105
Sample PPPoE server configuration using RADIUS	4-106
Administrative tools for IP routing.....	4-109

Chapter 5 Ethernet and IP QoS.....5-1

Overview of the QoS implementation	5-1
Packet classification subsystem	5-2
Rate-limiting, prioritization, and scheduling subsystem.....	5-3
Packet marking engine	5-3
Configuration steps.....	5-4
What the system does at the output interface	5-4
Default IP QoS configuration	5-5
Introduction to PACKET-FLOWS profile settings	5-5
Layer 2 classifiers	5-6
Packet classifiers.....	5-7
Caveat about fragmented IP packets	5-9
Details of packet classifier comparison passes	5-9
Comparisons of IP addresses	5-10
Comparisons of IP TOS values.....	5-10
Comparisons of port numbers.....	5-11
How nonmatching packets are prioritized (the default rule).....	5-11
Scheduling and rate limiting.....	5-12
Token buckets in the single-rate three color policing algorithm.....	5-13
Using a single rate two-color algorithm	5-14
Notes on the policing implementation	5-14
Rate limiting on VLAN bridged interfaces	5-14
Example of rate limiting on a BIR connection	5-15
Example of rate limiting on a 1:1 bridged VLAN	5-16
Packet marking	5-16
QoS packet marking for routed traffic.....	5-17
Ethernet p-bit marking for bridged or routed VLAN traffic	5-17
Overview of packet marking settings	5-18
Example of IP ToS marking on a routed VLAN interface.....	5-19

Example of Ethernet p-bit marking	5-20
Example of mapping ATM QoS to a packet marking value	5-21
QoS-related connection and interface settings.....	5-22
Applying a PACKET-FLOWS profile to an output interface	5-23
Inheritance of PACKET-FLOWS configurations on virtual IP interfaces ..	5-23
Virtual IP interfaces and interface grouping.....	5-24
ATM QoS and IP QoS considerations.....	5-24
Configuring Ethernet IP traffic shaping	5-24
Configuring VLAN Ethernet traffic shaping.....	5-25
QoS-related settings in the SYSTEM profile	5-25
Performance recommendations.....	5-25
Configurable queue size for IPTV traffic on DSL links.....	5-26
Tracking rate adaptation for downstream traffic shaping	5-26
How traffic shaping typically occurs	5-27
How rate adaptation affects traffic shaping.....	5-27
How to enable the rate-tracking optimization	5-27
How this feature works on the VDSL LIM	5-27
How this feature works on the ADSL2+ HB LIM.....	5-27
Examples of configuring QoS	5-28
Prioritizing IP packet flows based on DSL service contracts	5-29
Prioritizing different kinds of IP traffic on an ATM PVC.....	5-30
Prioritizing traffic using both IP and ATM QoS	5-32
Configuring bridging VLAN Ethernet QoS	5-36
Configuring traffic shaping and rate limiting for a N:1 VLAN bridging.....	5-37
Configuring rate limiting and traffic shaping with VLAN stacking.....	5-39
Configuring rate limiting and traffic shaping for 1:1 VLAN bridging	5-42
Configuring a PPPoE connection with and without line-rate tracking	5-44
Configuring a BIR connection with line-rate tracking	5-46
Administrative tools for monitoring IP QoS.....	5-47
Example of monitoring routed traffic onto Gigabit Ethernet	5-47
Creating a terminating routed connection	5-47
Applying a PACKET-FLOWS profile to the Ethernet IP interface.....	5-48
Obtaining the Ethernet interface number	5-48
Enabling monitoring on the Ethernet IP interface	5-48
Example of monitoring bridged VLAN traffic (transparent bridging).....	5-49
Creating a bridged VLAN interface and PACKET-FLOWS profile.....	5-49
Creating a bridged subscriber interface and PACKET-FLOWS profile	5-50
Obtaining the interface numbers	5-51
Enabling monitoring for the bridged VLAN interface	5-51
Enabling monitoring for the bridged subscriber interface.....	5-52
Example of monitoring 1:1 VLAN bridged traffic	5-53
Creating the subscriber-side profiles	5-53
Creating the VLAN-side profiles.....	5-54
Obtaining the interface numbers	5-54
Enabling monitoring in the upstream direction	5-55
Enabling monitoring in the downstream direction	5-56
Example of monitoring stacked VLAN bridged traffic	5-57
Enabling QoS monitoring on an NSP VLAN interface	5-57
Enabling QoS monitoring on a stacked VLAN WAN connection.....	5-58
Example of displaying a connection's traffic shaping parameters	5-59
Displaying a connection's traffic shaping parameters on the IP2000.....	5-59
Displaying a connection's traffic shaping parameters on the IP2100.....	5-60

Limitations with the current software version	5-61
Chapter 6 L2TP Tunneling Configuration	6-1
Overview of L2TP tunneling	6-2
Links to LNS servers	6-2
Links to PPP clients	6-2
L2TP-related network settings	6-3
L2TP-related RADIUS configuration	6-3
Overview of L2TP tunnel authentication	6-3
System name used for tunnel authentication	6-4
How the system finds a matching tunnel	6-4
Examples of how client-auth-id settings create parallel tunnels	6-5
Examples of configuration errors causing multiple tunnels	6-5
Configuring L2TP global options	6-6
Enabling L2TP operations and authentication	6-6
Setting L2TP timers and other variables	6-7
Retry timers	6-9
CSN encoding for RADIUS accounting	6-10
Sample global tunneling configuration	6-10
Configuring client connections with PPP authentication	6-11
Overview of mobile client connection settings	6-11
Sample PPP-authenticated client connection	6-13
Sample PPP-authenticated connection with two LNS systems	6-14
Configuring connection-based tunnel authentication	6-16
Overview of connection-based tunnel authentication settings	6-16
Example of connection-based tunnel authentication	6-17
Configuring server-based tunnel authentication	6-18
Overview of server-based tunnel authentication settings	6-19
Example TUNNEL-SERVER password configuration	6-20
Sample server-based authentication configuration	6-21
Using tunnel assignment IDs	6-22
Overview of tunnel assignment settings	6-22
Example of configuring tunnel assignment IDs	6-23
Administrative tools for monitoring L2TP	6-26
Displaying traffic statistics for open UDP/TCP ports	6-26
Displaying L2TP statistics	6-26
Displaying PPPoE statistics	6-27
Limitations with this software version	6-27
Impact of L2TP data sequencing limitation	6-28
Impact of L2TP IP QoS limitation	6-28
Chapter 7 Virtual Router Configuration	7-1
Overview of virtual routing	7-1
How virtual routers affect the routing table	7-2
Interconnecting virtual domains	7-2
Applicability and limitations	7-2
Creating a virtual router	7-3
Overview of VROUTER profile settings	7-3
Example of defining a virtual router	7-4
Defining address pools for a virtual router	7-7
Assigning interfaces to a virtual router	7-7

Overview of interface virtual router settings	7-7
Examples of assigning virtual router membership to interfaces	7-8
Defining virtual router static routes.....	7-8
Overview of static route settings	7-8
Examples of defining a route on a per-virtual-router basis	7-9
Specifying an inter-virtual-router route.....	7-10
Configuring virtual router DNS servers	7-10
Overview of virtual router DNS settings	7-11
Example of a typical virtual router DNS configuration	7-11
Deleting a virtual router.....	7-12
Administrative tools for virtual routers.....	7-12
 Chapter 8 OSPF Configuration	8-1
Overview of supported OSPF features	8-1
Limited border router capability	8-2
One active IP interface per port	8-2
Authentication	8-2
Support for variable-length subnet masks	8-2
Exchange of routing information	8-3
Designated and backup designated routers on broadcast networks	8-4
Routing across NBMA interfaces	8-5
Configurable cost metrics.....	8-5
Hierarchical routing (areas)	8-6
Link-state routing algorithms	8-7
Enabling OSPF systemwide	8-9
Configuring OSPF on Gigabit Ethernet	8-9
Overview of IP-INTERFACE OSPF settings.....	8-10
Sample Gigabit Ethernet interface configuration	8-12
Configuring OSPF on an ATM trunk interface	8-13
Overview of CONNECTION OSPF settings	8-13
Sample OSPF point-to-point configuration	8-14
Sample configuration of NBMA across point-to-point	8-14
Overview of additional NBMA settings.....	8-14
Example of an NBMA configuration	8-15
Configuring global route options that apply to OSPF	8-16
Example of importing a summarized pool as an ASE.....	8-17
Example of setting ASE preferences	8-17
Configuring IP-ROUTE OSPF options	8-17
Example of configuring a type 7 LSA in an NSSA.....	8-18
Example of assigning a cost to a static route	8-19
Administrative tools for OSPF routing.....	8-19
 Chapter 9 IP Multicast Configuration	9-1
IP multicast forwarding	9-1
Network-side MBONE interfaces.....	9-2
Notice about Gigabit Ethernet redundancy for a LAN MBONE	9-2
LIM-side multicast client interfaces.....	9-2
IP2100 and IP2000 maximum limits on multicast and bridge groups.....	9-3
Configuring MBONE interfaces.....	9-3
Overview of multiple MBONE configuration	9-4
Sample configuration with multiple MBONE interfaces	9-5

Sample MBONE configuration on Gigabit Ethernet VLANs.....	9-6
Managing multicast group memberships	9-8
Overview of MCAST-SERVICE settings.....	9-8
Sample multicast address filters.....	9-9
Sample multicast address range filter	9-10
Configuring multicast client interfaces.....	9-11
Overview of multicast client CONNECTION settings	9-11
Setting IGMP-v2 timers (local profiles only)	9-13
Example of using multiple multicast filters	9-14
Sample multicast video configuration with filters.....	9-15
Configuring the local MBONE interface.....	9-16
Configuring multicast client PVCs.....	9-16
Applying a filter that restricts the GigE interface to video traffic only	9-17
An alternative filter to restrict each client interface	9-18
Sample multicast video configuration with a remote MBONE interface.....	9-19
ATM QoS when both multicast and unicast clients are supported.....	9-21
Multicast server virtual circuits.....	9-22
Overview of multicast server VC settings	9-23
Sample configuration of multicast server VCs	9-24
Enabling the multicast server VC feature.....	9-24
Configuring the multicast server VCs.....	9-24
Configuring the terminating connections for DSL users.....	9-25
Distributed multicast with HB LIMs (IP2100 only).....	9-25
Benefits of distributed multicast	9-25
Conceptual overview	9-26
MBONE and multicast client interactions on FFW interfaces	9-27
What happens when a CPE joins a multicast group	9-27
What happens when a CPE leaves a multicast group	9-28
Multicast-capable HB LIM considerations	9-28
Downstream traffic processing	9-28
Changes in packet-flow prioritization with distributed multicast.....	9-29
Scheduling priority of the HB LIM FFW interface	9-30
ATM QoS for unicast traffic on the HB LIM.....	9-31
Statistics and diagnostics for the multicast-capable HB LIM.....	9-31
Administrative tools for IGMP operations.....	9-31
CAC for multicast channels.....	9-31
MULTICAST-CAC	9-32
Multicast CAC specific parameters in the SYSTEM profile	9-35
Future Provisions	9-36
Configuring multicast CAC on Stinger CM.....	9-38
Viewing multicast CAC statistics	9-39
Viewing CAC statistics using CLI.....	9-39
Viewing CAC statistics using SNMP	9-39

Chapter 10 PIM-SM v2 Configuration 10-1

PIM-SM features supported with this software version.....	10-2
Overview of PIM-SM configuration	10-2
Enabling multicast and PIM.....	10-3
Overview of settings in the IP-GLOBAL profile	10-3
Example showing BSR election and dynamic group-RP mappings	10-4
Configuring static mappings between groups and rendezvous points.....	10-5

Configuring PIM on Gigabit Ethernet or trunk interfaces	10-6
PIM options in the IP-INTERFACE and CONNECTION profiles	10-7
Example of enabling PIM on the Gigabit Ethernet interface	10-9
Example of enabling PIM on a trunk interface	10-10
Sample PIM-SM system configuration	10-10
Administrative tools for PIM-SM routing	10-11
Chapter 11 Filter Configuration	11-1
Filter overview	11-1
Filter rules	11-2
Explicit default filter rules	11-2
Defining IP filters	11-2
Overview of IP filter settings	11-2
Details of IP filter comparison passes	11-4
Filtering on source or destination IP addresses	11-5
Filtering on port numbers	11-5
TCP-established filters	11-6
Sample IP filters	11-7
Sample TCP-established filter	11-8
Preventing address spoofing	11-8
An IP filter for more complex security issues	11-9
Sample filter with no explicit default rule	11-10
Sample filter with explicit default rule	11-11
Sample filter using a generic explicit default rule	11-11
Defining ICMP filters	11-12
Overview of profile settings	11-12
Sample ICMP filter configurations	11-14
Blocking incoming Echo-Request packets	11-14
Stopping Echo-Request packets from being forwarded upstream	11-16
Blocking Echo-Requests to a specific IP address	11-16
Defining route filters	11-17
Overview of route filter settings	11-17
Sample route filters	11-18
Sample route filter that excludes a route	11-18
Sample route filter that configures a route's metric	11-19
Defining Ethernet input filters	11-19
Overview of ethernet filter settings	11-19
Sample PPPoE and MAC address filter	11-20
Applying a filter to an interface	11-21
Profile settings for applying a filter	11-21
Applying filters to a CPE interface	11-21
Applying a filter to an Ethernet interface	11-22
Applying a filter to a VLAN Ethernet interface	11-22
Applying an Ethernet filter to a VLAN bridging connection	11-23
Administrative tools for filters	11-23
Appendix A IP Control Module Diagnostics	A-1
Enabling the debug environment	A-2
Gigabit Ethernet diagnostics	A-2
IGMP diagnostics	A-4
PIM-SM diagnostics	A-10

VLAN-related diagnostics	A-13
SAR-related diagnostics.....	A-19
Network processor-related diagnostics.....	A-20
Multicast-capable high-bandwidth LIM diagnostics	A-27
SNMP MIB for GMAC and VLAN statistics	A-29
History maintained at 15-minute intervals.....	A-29
Gigabit Ethernet (GigE) statistics tables	A-29
Gigabit Ethernet configuration.....	A-30
Interval transmit statistics	A-30
Total transmit statistics.....	A-30
Interval receive statistics	A-32
Total receive statistics.....	A-32
VLAN statistics tables	A-34
VLAN statistics.....	A-34
VLAN clear statistics	A-35
PIMv2 MIB support.....	A-35
Index	Index-1

Figures




Figure 1-1	IP2100 control module basic architecture.....	1-6
Figure 1-2	Sample setup showing multicast and unicast video services	1-16
Figure 1-3	Sample setup showing Internet access and voice over ATM	1-16
Figure 1-4	Sample setup showing multiple IP flows to a CPE router	1-17
Figure 2-1	Gigabit Ethernet redundancy for RFC 2684 connectivity	2-6
Figure 2-2	Gigabit Ethernet redundancy for a LAN MBONE	2-8
Figure 2-3	Redundant uplinks for a VLAN.....	2-18
Figure 3-1	Bridging VLAN: One PVC to one VLAN (1:1)	3-2
Figure 3-2	Sample 1:1 VLAN circuit	3-5
Figure 3-3	Bridging multiple PVCs to a VLAN.....	3-6
Figure 3-4	Sample N:1 VLAN bridging	3-12
Figure 3-5	Format when an interface IP address (if-ip) is used	3-15
Figure 3-6	Format when a text string (vendor-option-string) is used	3-15
Figure 3-7	Format when sending only the hostname (send-only-hostname) ..	3-15
Figure 3-8	DHCP snooping example.....	3-17
Figure 3-9	Contents of fields with sample DHCP snooping configuration	3-19
Figure 3-10	PPPoE vendor-specific tag format	3-21
Figure 3-11	PPPoE snooping example	3-23
Figure 3-12	Contents of fields with sample PPPoE snooping configuration	3-24
Figure 3-13	Stacked VLAN: Bridging untagged frames from DSL interfaces.....	3-28
Figure 3-14	Stacked VLAN: Bridging enterprise VLAN-tagged frames.....	3-31
Figure 3-15	Sample routed VLAN	3-35
Figure 3-16	Sample routed VLAN	3-36
Figure 3-17	Terminating IPoA traffic on an upstream BRAS	3-39
Figure 3-18	Sample 1:1 VLAN circuits for bridging IPoA	3-42
Figure 3-19	Two BRAS with the same IP address in different VLANs.....	3-44
Figure 3-20	Not recommended with unnumbered interfaces	3-46
Figure 3-21	Sample transparent-bridged PPPoA configuration	3-50
Figure 3-22	Sample PVC-to-VLAN bridged PPPoA configuration	3-53
Figure 4-1	Four DSL links using ECMP routes	4-17
Figure 4-2	Two routed VLAN uplinks using ECMP routes	4-21
Figure 4-3	Remote CPE requiring dynamic IP address assignment	4-29
Figure 4-4	Example of a PPPoE session using slot-based address assignment ..	4-37
Figure 4-5	Default route to a local IP router.....	4-41
Figure 4-6	Static route to a subnet	4-42
Figure 4-7	Sample configuration using two soft IP interface addresses.....	4-45
Figure 4-8	Sample network with two levels of anti-spoofing protection	4-48
Figure 4-9	Router-to-router IP connection	4-51
Figure 4-10	A numbered-interface connection	4-52
Figure 4-11	Forwarding terminating PVCs on the Gigabit Ethernet interface ...	4-53
Figure 4-12	Aggregating PVCs onto a single virtual circuit using IP routing.....	4-55

Figure 4-13	BIR interface on a LIM port	4-56
Figure 4-14	BIR subnet configuration on LIM interface	4-57
Figure 4-15	BIR/32 configurations	4-59
Figure 4-16	Bidirectional filtering on a BIR interface.....	4-61
Figure 4-17	Sample BIR connection with four static IP addresses	4-64
Figure 4-18	DHCP relay for an IPoA terminated PVC	4-74
Figure 4-19	Option 82 information field formats	4-76
Figure 4-20	Sample DHCP usage with LAN management interface.....	4-77
Figure 4-21	Sample DHCP setup on BIR connection	4-78
Figure 4-22	PPPoA topology	4-86
Figure 4-23	PPPoE topology	4-87
Figure 4-24	Example of a PPPoA session on a DSL interface	4-91
Figure 4-25	Example of a PPPoE session on a DSL interface.....	4-95
Figure 5-1	QoS subsystems.....	5-1
Figure 5-2	Sample deployment with VLAN p-bit marking.....	5-17
Figure 5-3	Using interface grouping to prioritize traffic by service level.....	5-29
Figure 5-4	Unicast and multicast video share the same priority	5-31
Figure 5-5	Prioritizing traffic at the connection level and flow level	5-32
Figure 5-6	Stacked VLAN requiring p-bit remarking on the output interface .	5-36
Figure 5-7	Sample N:1 VLAN bridging with traffic shaping and rate limiting..	5-37
Figure 5-8	Stacked VLAN with rate limiting and traffic shaping.....	5-39
Figure 5-9	1:1 bridging VLAN with rate limiting and traffic shaping	5-42
Figure 6-1	L2TP tunneling.....	6-2
Figure 6-2	Bringing up a tunnel using PPP authentication	6-13
Figure 6-3	Primary and secondary tunnel end points	6-15
Figure 6-4	Connection-based tunnel authentication.....	6-17
Figure 6-5	Server-based tunnel authentication	6-21
Figure 6-6	Tunnel assignment IDs.....	6-23
Figure 7-1	Simple diagram of three virtual domains (virtual routers)	7-1
Figure 8-1	OSPF broadcast network on Gigabit Ethernet.....	8-4
Figure 8-2	OSPF costs for different types of links.....	8-5
Figure 8-3	Dividing an OSPF autonomous system into areas	8-6
Figure 8-4	Sample OSPF topology	8-7
Figure 8-5	OSPF on a LAN interface.....	8-12
Figure 8-6	OSPF over ATM point to point.....	8-14
Figure 8-7	OSPF NBMA over ATM point to point.....	8-15
Figure 9-1	Multicast video sample setup	9-1
Figure 9-2	Multiple MBONE interfaces on trunk or LAN interfaces	9-2
Figure 9-3	Sample configuration with multiple MBONE interfaces.....	9-5
Figure 9-4	Sample configuration of VLAN MBONE interface	9-7
Figure 9-5	DSL video application with a local MBONE interface.....	9-16
Figure 9-6	IPTV video sample configuration	9-19
Figure 9-7	One VC per multicast group for incoming multicast data streams..	9-22
Figure 9-8	Multicast server VCs on a trunk interface.....	9-24
Figure 9-9	FFW interfaces on multicast-capable modules.....	9-26
Figure 9-10	Two IGMP Join requests for the same multicast group	9-27
Figure 9-11	Two IGMP Leave requests for the same multicast group.....	9-28
Figure 10-1	PIM-SM on Gigabit Ethernet and trunk interface.....	10-10

Tables




Table 1-1	Overview of network features.....	1-1
Table 1-2	IP2000 model numbers and platform support	1-3
Table 1-3	IP2100 control module enhancements	1-5
Table 1-4	IP2100 model numbers and platform support	1-6
Table 1-5	NP port manager profiles	1-9
Table 1-6	Default classification and prioritization	1-17
Table 3-1	Definition of VLAN terms.....	3-1
Table 3-2	Traffic restrictions when port blocking is enabled.....	3-9
Table 3-3	IGMP control packet handling with IGMP snooping enabled.....	3-10
Table 3-4	Maximum number of learned MAC addresses per module.....	3-11
Table 3-5	Packet handling with PPPoE snooping.....	3-20
Table 3-6	Enterprise VLAN tagged frames mapped to a stacked VLAN	3-33
Table 3-7	ARP request packet contents on a configured VLAN	3-48
Table 3-8	VLAN ID argument added to arptable command.....	3-56
Table 4-1	Decimal subnet masks and corresponding prefix lengths	4-5
Table 4-2	Profiles for configuring logical IP interfaces for Ethernet ports.....	4-6
Table 4-3	Settings that affect contents of the giaddr field in DHCP packets ...	4-74
Table 4-4	Required encapsulation types for PPPoA and PPPoE	4-90
Table 4-5	Configurable LCP Echo interaction with LQM.....	4-100
Table 4-6	Permanent RADIUS PPPoE server profiles.....	4-107
Table 5-1	Descriptions of QoS support for routed, bridged, and VLAN traffic..	5-2
Table 5-2	Packet marking supported on egress interfaces.....	5-4
Table 5-3	Comparison passes performed on inbound packet flows	5-9
Table 5-4	Rate limiting terminology	5-12
Table 5-5	Flow-based rate limiting configurations.....	5-14
Table 5-6	VLAN traffic shaping configurations	5-25
Table 5-7	Sample IP traffic types and priorities across an ATM PVC	5-31
Table 5-8	Sample IP traffic types and priorities across two PVCs.....	5-32
Table 5-9	Sample bandwidth limitations	5-33
Table 5-10	Ethernet p-bit remarking table.....	5-36
Table 6-1	Network settings related to L2TP	6-3
Table 6-2	Existing tunnels to the same LNS	6-5
Table 6-3	Tunnels created for clients based on profile settings.....	6-5
Table 6-4	Tunnels created when user1 dials in first (error not detected)	6-6
Table 6-5	Tunnels created when user2 dials in first (error shown)	6-6
Table 7-1	Administrative commands showing optional vrouter arguments...	7-12
Table 8-1	Description of LSA types	8-3
Table 8-2	Link-state databases for OSPF topology in Figure 8-4	8-7
Table 8-3	Shortest-path tree and resulting routing table for Router-1	8-8
Table 8-4	Shortest-path tree and resulting routing table for Router-2	8-8
Table 8-5	Shortest-path tree and resulting routing table for Router-3	8-8
Table 9-1	Maximum multicast groups, multicast users, and bridge groups.....	9-3

Tables

Table 9-2	Sample multicast and unicast client requirements	9-21
Table 9-3	FFW interfaces registered as MBONE or multicast client.....	9-27
Table 9-4	Packet flow with multicast traffic as the highest priority flow	9-29
Table 9-5	Packet flow with unicast traffic as the highest priority flow	9-29
Table 9-6	Packet flow with interleaving priorities (unicast/multicast)	9-30
Table 9-7	Packet flow with interleaving priorities (multicast/unicast)	9-30
Table 10-1	Current level of support for PIM-SM functionality.....	10-2
Table 11-1	Default filtering behavior	11-1
Table A-1	GigEConfigTable MIB objects.....	A-30
Table A-2	GigETxIntervalTable MIB objects.....	A-30
Table A-3	GigETxTotalTable MIB objects	A-31
Table A-4	GigERxIntervalTable MIB objects	A-32
Table A-5	GigERxTotalTable MIB objects	A-33
Table A-6	GigEVlanStatTable MIB objects.....	A-34
Table A-7	GigEVlanClearStatTable MIB objects	A-35
Table A-8	Current level of support for PIMv2 MIB tables.....	A-35

About This Guide

A Stinger system with an IP control module supports Asynchronous Transfer Mode (ATM) capabilities similar to those in Stinger units with a standard control module, as described in the *Stinger ATM Configuration Guide*. In addition, an IP control module can terminate and route IP and related traffic.



Note Instructions for installing and configuring the management functions of the Stinger system are found in the *Getting Started Guide* for your Stinger platform.

What is in this guide

This guide focuses on the aspects of Stinger configuration that are specific to the IP DSLAM functionality. It describes how to configure IP routing and related functions in the Stinger. It includes information about local and global network IP issues, as well as how to configure both IP-routed switch-through ATM permanent virtual circuits (PVCs) and RFC 2684 PVCs.

This guide also describes how to set up IEEE 802.1Q virtual local area network (VLAN) support on the Gigabit Ethernet interface, and how to configure the system to support multicast video over DSL with Internet Group Management Protocol (IGMP) version-1 or version-2 messaging.

You can configure the amount of bandwidth allocated to LIM interfaces and control modules for carrying upstream traffic. For details about that aspect of using the IP control module, as well as for ATM PVCs, ATM quality of service (QoS) and other traffic management capabilities, see the *Stinger ATM Configuration Guide*.



Warning Before installing your Stinger unit, be sure to read the safety instructions in the *Edge Access and Broadband Access Safety and Compliance Guide*. For information specific to your unit, see the “Safety-Related Physical, Environmental, and Electrical Information” appendix in the *Getting Started Guide* for your Stinger unit.

Documentation conventions

Following are all the special characters and typographical conventions used in this manual:

Convention	Meaning
Monospace text	Represents text that appears on your computer's screen, or that could appear on your computer's screen.
Boldface monospace text	Represents characters that you enter exactly as shown (unless the characters are also in <i>italics</i> —see <i>Italics</i> , below). If you could enter the characters but are not specifically instructed to, they do not appear in boldface.
<i>Italics</i>	Represent variable information. Do not enter the words themselves in the command. Enter the information they represent. In ordinary text, italics are used for titles of publications, for some terms that would otherwise be in quotation marks, and to show emphasis.
[]	Square brackets indicate an optional argument you might add to a command. To include such an argument, type only the information inside the brackets. Do not type the brackets unless they appear in boldface.
	Separates command choices that are mutually exclusive.
>	Points to the next level in the path to a parameter or menu item. The item that follows the angle bracket is one of the options that appear when you select the item that precedes the angle bracket.
Key1+Key2	Represents a combination keystroke. To enter a combination keystroke, press the first key and hold it down while you press one or more other keys. Release all the keys at the same time. (For example, Ctrl+H means hold down the Ctrl key and press the H key.)
Press Enter	Means press the Enter or Return key or its equivalent on your computer.
 Note	Introduces important additional information.
 Caution	Warns that a failure to follow the recommended procedure could result in loss of data or damage to equipment.
 Warning	Warns that a failure to take appropriate safety precautions could result in physical injury.
 Warning	Warns of danger of electric shock.

Acronyms used in this guide

AAA	Authentication, Authorization, and Accounting
AAL5	Asynchronous Transfer Mode Adaptation Layer 5
ABR	Area Border Router
ARP	Address Resolution Protocol
AS	Autonomous System
ASBR	Autonomous System Border Router
ASE	Autonomous System External
ATM	Asynchronous Transfer Mode
BDR	Backup Designated Router
BIR	Bridged IP Routing
BOOTP	Bootstrap Protocol
BRAS	Broadband Remote Access Server
BSM	Bootstrap Message
BSR	Bootstrap Router
C-RP	Candidate Rendezvous Point
CAC	Connection Admission Control
CBS	Committed Burst Size
CBR	Constant Bit Rate
CDMA/CD	Code Division Multiple Access/Carrier Detect
CHAP	Challenge Handshake Authentication Protocol
CIR	Committed Information Rate
CoS	Class of Service
CPE	Customer Premises Equipment
CRC	Cyclic Redundancy Check
DHCP	Dynamic Host Configuration Protocol
DID	Destination ID
DNS	Domain Name System
DSL	Digital Subscriber Loop
DSLAM	DSL Access Multiplexer
DR	Designated Router
DVRMP	Distance Vector Multicast Routing Protocol
EBS	Excess Burst Size
ECMP	Equal Cost Multi Path
EMI	Electromagnetic Interference
GE-OLIM	Gigabit Ethernet Optical Line Interface Module
GMAC	Gigabit Ethernet Media Access Control
HB	High Bandwidth

IAD	Integrated Access Device
ICMP	Internet Control Message Protocol
IGMP	Internet Group Management Protocol
IGP	Interior Gateway Protocol
IP	Internet Protocol
IPCP	IP Control Protocol
IPoA	Internet Protocol over ATM
ISP	Internet Service Provider
L2TP	Layer 2 Tunneling Protocol
LACP	Link Aggregation Control Protocol
LAG	Link Aggregation Group
LAMP	Link Aggregation Marker Protocol (not yet supported)
LAN	Local Area Network
LCP	Link Control Protocol
LIM	Line Interface Module
LLC	Logical Link Control
LSA	Link State Advertisement
MAC	Media Access Control
MBONE	Multicast Backbone
MIB	Management Information Base
MOSPF	Multicast OSPF
MPOA	Multiprotocol over ATM
MP	Multilink Protocol
MTU	Maximum Transmission Unit
NBMA	Non-Broadcast Multi Access
NP	Network Processor
NSP	Network Service Provider
NSSA	Not So Stubby Area
OAM	Operations, Administration, and Maintenance
OSPF	Open Shortest Path First
PADI	PPPoE Active Discovery Initiation
PADO	PPPoE Active Discovery Offer
PADR	PPPoE Active Discovery Request
PADS	PPPoE Active Discovery Session-confirmation
PADT	PPPoE Active Discovery Terminate
PAP	Password Authentication Protocol
PCR	Peak Cell Rate
PDU	Protocol Data Unit
PIM-SM	Protocol Independent Multicast-Sparse Mode

PMBR	PIM Multicast Border Router
PNNI	Private Network-to-Network Interface
POP	Point of Presence
POST	Power-On Self Test
PPP	Point-to-Point Protocol
PPPoA	Point-to-Point over ATM
PPPoE	Point-to-Point over Ethernet
PPV	Pay Per View
PVC	Permanent Virtual Circuit
QoS	Quality of Service
RADIPAD	RADIUS IP Address Daemon
RADIUS	Remote Authentication Dial In User Service
RARP	Reverse Address Resolution Protocol
RAS	Remote Access Server
RIP	Routing Information Protocol
RFC	Request for comment
RLIM	Remote Line Interface Module
RP	Rendezvous Point
RPF	Reverse Path Forwarding
RPT	Rendezvous Point Tree
RSTP	Rapid Spanning Tree Protocol
SAR	Segmentation and Reassembly
SCCRP	Start Control Connection Reply
SCCRQ	Start Control Connection Request
SCR	Sustainable Cell Rate
SFP	Small Form Factor Pluggable transceiver
SPT	Shortest Path Tree
STB	Set Top Box
STP	Spanning Tree Protocol
TACL	Telnet Access Control Lists
TOS	Type of Service
UBR	Unspecified Bit Rate
UDP	User Datagram Protocol
UNI	User Network Interface
VBR	Variable Bit Rate
VC	Virtual Circuit
VCC	Virtual Channel Connection
VCI	Virtual Channel Identifier
VLAN	Virtual Local Area Network

VLSM	Variable Length Subnet Mask
VP	Virtual Path
VPC	Virtual Path Connection
VPI	Virtual Path Identifier
VPN	Virtual Private Network
VSA	Vendor-Specific Attributes
WAN	Wide Area Network

Stinger documentation set

The Stinger documentation set consists of the following manuals, which can be found at <http://www.lucent.com/support>.

■ Read me first:

- *Edge Access and Broadband Access Safety and Compliance Guide*. Contains important safety instructions and country-specific information that you must read before installing a Stinger unit.
- *TAOS Command-Line Interface Guide*. Introduces the TAOS command-line environment and shows you how to use the command-line interface effectively. This guide describes keyboard shortcuts and introduces commands, security levels, profile structure, and parameter types.

■ Installation and basic configuration:

- *Getting Started Guide* for your Stinger platform. Shows how to install your Stinger chassis and hardware. This guide also shows you how to use the command-line interface to configure and verify IP access and basic access security on the unit, and how to configure Stinger control module redundancy on units that support it.
- *Stinger Compact Remote Installation and Configuration Guide*. Provides an overview of the Stinger Compact Remote and instructions for the installation and replacement of its components. This guide also describes how to configure and manage the Compact Remote as a hosted unit.
- Module guides. For each Stinger line interface module (LIM), trunk module, or other type of module, an individual guide describes the module's features and provides instructions for configuring the module and verifying its status.

■ Configuration:

- *Stinger ATM Configuration Guide*. Describes how to integrate the Stinger into the ATM and Digital Subscriber Line (DSL) access infrastructure. The guide explains how to configure PVCs, and shows how to use standard ATM features such as quality of service (QoS), connection admission control (CAC), and subtending.
- *Stinger IP Control Module Configuration Guide*. For Stinger systems with an IP control module, this guide describes how to integrate the system into the IP infrastructure. Topics include IP-routed switch-through ATM PVCs and RFC 2684 PVCs, IEEE 802.1Q VLAN, and forwarding multicast video transmissions on DSL interfaces.
- *Stinger Private Network-to-Network Interface (PNNI) Supplement*. For the optional PNNI software, this guide provides quick-start instructions for configuring

- PNNI and soft PVCs (SPVCs), and describes the related profiles and commands.
- *Stinger SNMP Management of the ATM Stack Supplement*. Describes SNMP management of ATM ports, interfaces, and connections on a Stinger unit to provide guidelines for configuring and managing ATM circuits through any SNMP management utility.
- *Stinger T1000 Module Routing and Tunneling Supplement*. For the optional T1000 module, this guide describes how to configure the Layer 3 routing and virtual private network (VPN) capabilities.
- **RADIUS: TAOS RADIUS Guide and Reference**. Describes how to set up a unit to use the Remote Authentication Dial-In User Service (RADIUS) server and contains a complete reference to RADIUS attributes.
- **Administration and troubleshooting: Stinger Administration Guide**. Describes how to administer the Stinger unit and manage its operations. Each chapter focuses on a particular aspect of Stinger administration and operations. The chapters describe tools for system management, network management, and Simple Network Management Protocol (SNMP) management.
- **Reference:**
 - *Stinger Reference*. An alphabetic reference to Stinger profiles, parameters, and commands.
 - *TAOS Glossary*. Defines terms used in documentation for Stinger units.

Related documents

The following industry documents provide background information about features described in this guide:

- RFC 951, *Bootstrap Protocol*
- RFC 1112, *Host Extensions for IP Multicasting*
- RFC 1334, *PPP Authentication Protocols*
- RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*
- RFC 1587, *The OSPF NSSA Option*.
- RFC 1700, *Assigned Numbers*
- RFC 1723, *RIP Version 2: Carrying Additional Information*
- RFC 1812, *Requirements for IP Version 4 Routers*
- RFC 1994, *PPP Challenge Handshake Authentication Protocol (CHAP)*
- RFC 2131, *Dynamic Host Configuration Protocol*
- RFC 2132, *DHCP Options and BOOTP Vendor Extensions*
- RFC 2236, *Internet Group Management Protocol Version 2*
- RFC 2328, *OSPF Version 2*
- RFC 2362, *Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification*, draft-ietf-pim-sm-v2-new-07.txt, March 2003, draft-ietf-pim-sm-bsr-03.txt, February 2003
- RFC 2364, *PPP over AAL5*
- RFC 2516, *A Method for Transmitting PPP Over Ethernet (PPPoE)*

- RFC 3046, *DHCP Relay Agent Information Option*
- RFC 2697, *A Single Rate Three Color Marker*
- IEEE 802.1Q-1998, *IEEE Standard for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks*
- IEEE 802.1P, *LAN Layer 2 QoS/CoS Protocol for Traffic Prioritization*

Introduction

1

Stinger IP DSLAM network features	1-1
Introduction to the Stinger IP2000 control module.	1-3
Introduction to the Stinger IP2100 control module.	1-5
Network architecture overview	1-15

A Stinger system with an IP control module provides cost-effective, high-speed DSL access using native IP architecture. A Stinger system with an IP control module is referred to in this guide as a *Stinger IP DSLAM*.

Stinger IP DSLAM network features

A Stinger IP control module supports both an IP network processor and ATM switching fabric that supports the same set of features as non-IP-capable control modules. For details about ATM features, see the *Stinger ATM Configuration Guide*.

Table 1-1 provides a general overview of the IP network-related features supported by Stinger IP control modules:

Table 1-1. Overview of network features

Category	Features
Layer 3 Routing and QoS	RIPv1, RIPv2, OSPF routing protocols
	DHCP relay with option 82
	IP filters, route filters
	Virtual routers
	IP QoS per IP flow with 8 priority queues
	TOS bit packet marking
	Rate limiting
	IPoA (RFC 2684), BIR connections
	Equal cost multipath (ECMP) load balancing (IP2000 and IP2100)

Table 1-1. Overview of network features (Continued)

Category	Features
Multicasting	Multicast protocols: PIM-SM
	IGMP forwarding (v1, v2)
	Multicast group management
	Multicast filters
Layer 2 Bridging and QoS	PVC-to-VLAN bridging (1:1)
	VLAN stacking
	Transparent bridging (IEEE 802.1d) (N:1 VLAN)
	Classification and Priority-bit marking
	IGMP snooping
	DHCP snooping with option 82
	PPPoE intermediate agent
	Ethernet filters
	VLAN traffic shaping
	Forwarding traffic to a VLAN based on Class of traffic
	Link Aggregation Control Protocol (LACP) on IP2100 only
	Rapid Spanning Tree Protocol (RSTP) on IP2100 only
	Bridging IPoA and PPPoA traffic to an upstream BRAS
Broadband RAS	Subscriber termination
	PPPoA/PPPoE session termination
	PPPoA/PPPoE autodetection capability
	L2TP tunneling support

Table 1-1. Overview of network features (Continued)

Category	Features
Security	IP anti-spoofing
	Device attachment limiting per interface
	Telnet access control lists
	SSH-v2
	RADIUS, extended RADIUS
	MS-CHAP
	User groups and password expiration
	Command logging per user profile
	ICMP filters
	Password Authentication Protocol (PAP)
	Challenge Authentication Protocol (CHAP)
	Profile-based access

Introduction to the Stinger IP2000 control module

The Stinger IP2000 control module provides one fiber-based Gigabit Ethernet interface, which uses modular Small Form Factor Pluggable (SFP) transceivers, and a network processor (NP) capacity of approximately 2.5Gbps.

The IP2000 also supports one 10/100 BASE-T Ethernet port, and one RS-232 serial port, for management access.

IP2000 model numbers and platform support

Table 1-2 shows the IP2000 model numbers and platform support:

Table 1-2. IP2000 model numbers and platform support

Control module number	Description	Supporting platforms
STGR-CM-IP2000-F	IP2000 control module	Stinger FS/Stinger FS+/ Stinger LS Stinger Compact Remote
STGRRT-CM-IP2000-F	IP2000 control module environmentally hardened	Stinger RT Stinger FS/Stinger FS+/ Stinger LS Stinger Compact Remote
STGR-SFP-SX	Short-haul SFP module	Stinger FS/Stinger FS+/ Stinger LS
STGR-SFP-LX	Long-haul SFP module	Stinger FS/Stinger FS+/ Stinger LS

Table 1-2. IP2000 model numbers and platform support (Continued)

Control module number	Description	Supporting platforms
STGRRT-SFP-LX	Long-haul SFP module, environmentally hardened	Stinger RT Stinger Compact Remote

IP2000 support for up to 2048 trunk terminated calls

Stinger systems with an IP2000 control module can establish 1175 terminated calls with default settings, and with the proper configuration can establish 2048 terminated calls. To enable the higher number of terminated calls, you must reduce the number of VCIs supported on at least one of the system's LIM slots. On a fully loaded system, reducing the number of VCIs on at least six LIM slots is recommended.

Following is the parameter, shown with its default value for a LIM in slot 1, for restricting the number of VCIs supported on the LIM slot:

```
[in SLOT-STATIC-CONFIG/{ shelf-1 slot-1 0 }]
num-of-vci = two-k
```

This parameter applies only to the LIM slots of a standalone or host Stinger system. A system reset is required to recognize the new value.



Note Changing the default setting for this parameter is not recommended for OLIMs on a fully loaded Stinger system.

Parameter

num-of-vci

Setting

Maximum number of VCIs that can be allocated on the module in this slot. Valid values are **one-k** (to indicate that 1024 VCIs can be allocated by the card) and **two-k** (the default, which indicates that 2048 VCIs are available for allocation).

When all SLOT-STATIC-CONFIG profiles use the default **two-k** setting, the system can bring up 1175 terminated calls. If one or more SLOT-STATIC-CONFIG profile is modified to a value of **one-k**, an increased number of VCIs are made available to the system's control module, which enables it to terminate an increased number of calls, up to a maximum of 2048. On a fully loaded system, reducing the number of VCIs on at least six LIM slots is recommended.

When you write the SLOT-STATIC-CONFIG profile after modifying the num-of-vci setting, the system displays a warning message indicating that you must reset the system for the new setting to take effect. For example:

```
admin> read slot-static-c { 1 1 0 }
admin> set num-of-vci = one-k
admin> write -f
```



```
LOG warning, Shelf 1, Controller-1, Time: 01:08:09--
    Please reset the system immediately to make this new change effective and
    avoid any undesired behavior
admin> reset
```

Features not currently supported by the IP2000

The IP2000 control module does not currently support the following features:

- Link Aggregation Control Protocol (LACP)
- Rapid Spanning Tree Protocol (RSTP)
- IP filters for Ethernet-encapsulated bridged IP datagrams
- Distributed multicast

Introduction to the Stinger IP2100 control module

The IP2100 control module is the successor to the IP2000 control module. It provides the same basic set of capabilities plus the following enhancements:

Table 1-3. IP2100 control module enhancements

Feature	Description
Gigabit Ethernet	Dual Gigabit Ethernet uplink ports
Network processor	4Gbps network processor
Backplane speed	Supports STS-12c (622Mbps) backplane links to high-bandwidth (HB) Annex A LIMs with ADSL2+ capability (STGR-LIM-AP-72-HB and STGR-LIM-AP-48-HB)
	Supports STS-3c (155Mbps) backplane links to legacy xDSL and T1/E1 LIMs
	Supports configurable network processor port managers to distribute NP capacity
Operating temperature	Designed for industrial temperature range of -40C to +65C

The IP2100 supports up to 4,096 trunk terminating connections, as compared to 2,048 for the IP2000.


With its increased throughput and support for high-bandwidth (HB) Annex A LIMs with ADSL2+ capability (STGR-LIM-AP-72-HB and STGR-LIM-AP-48-HB), the IP2100 provides performance enhancements useful for high-quality IP video streaming.

Introduction

Introduction to the Stinger IP2100 control module

Figure 1-1 shows the basic system architecture using the IP2100:

Figure 1-1. IP2100 control module basic architecture



When you install an HB LIM, the IP2100 generates a log message to the effect that you must modify the SLOT-STATIC-CONFIG profile for that slot to allocate the additional bandwidth. For details about bandwidth allocation, see the *Stinger ATM Configuration Guide*.

Model numbers and platform support

Table 1-4 shows control module model numbers and platform support:

Table 1-4. IP2100 model numbers and platform support

Control module number	Description	Supporting platforms
STGR-CM-IP2100-F-MODEL	IP2100 control module	Stinger FS/Stinger FS+/ Stinger LS Stinger Compact Remote
STGR-SFP-SX	Short-haul SFP module	Stinger FS/Stinger FS+/ Stinger LS Stinger Compact Remote
STGR-SFP-LX	Long-haul SFP module	Stinger FS/Stinger FS+/ Stinger LS Stinger Compact Remote

IP2100 memory modes

The IP2100 control module supports selectable operation modes for its fast pattern processor (FPP) classification program memory. By default, it uses a high-performance operation mode, which maximizes classification speed and uses 8MB memory for classification. In a capacity environment, you could see log messages indicating that FPP memory is low and calls could be dropped as a result. In that case, the number of calls the system can simultaneously process can be pushed higher by changing the mode.

Following is the relevant parameter, shown with its default setting:

```
[in SYSTEM]
np-fpp-class-mem-mode = high-performance
```

Parameter	Setting						
np-fpp-class-mem-mode	Specifies the memory mode for operating the IP2100 control module FPP classification program memory. If you change this value, you must reset the system for the new value to become operational. See “Required IP2100 memory mode to maximize user connections” on page 1-8 for related information. Following are valid values: <table><tr><td>high-performance</td><td>This is the default mode and the most efficient for accessing program memory. With this mode, classification is fast, but less memory (approximately 8MB) is available for classification.</td></tr><tr><td>high-capacity-0</td><td>This mode is less efficient in accessing program memory but provides an increase in the amount of program memory. With this mode, classification occurs more slowly but more memory (approximately 32MB) is available for classification.</td></tr><tr><td>medium-capacity</td><td>With this mode, classification occurs at a rate between the high-performance and high-capacity settings, and available memory for classification is approximately 16MB.</td></tr></table> This setting is required to maximize the number of user connections in the system.	high-performance	This is the default mode and the most efficient for accessing program memory. With this mode, classification is fast, but less memory (approximately 8MB) is available for classification.	high-capacity-0	This mode is less efficient in accessing program memory but provides an increase in the amount of program memory. With this mode, classification occurs more slowly but more memory (approximately 32MB) is available for classification.	medium-capacity	With this mode, classification occurs at a rate between the high-performance and high-capacity settings, and available memory for classification is approximately 16MB.
high-performance	This is the default mode and the most efficient for accessing program memory. With this mode, classification is fast, but less memory (approximately 8MB) is available for classification.						
high-capacity-0	This mode is less efficient in accessing program memory but provides an increase in the amount of program memory. With this mode, classification occurs more slowly but more memory (approximately 32MB) is available for classification.						
medium-capacity	With this mode, classification occurs at a rate between the high-performance and high-capacity settings, and available memory for classification is approximately 16MB.						

Maximizing the number of user connections

An IP2100 system creates internal OAM (operations, administration, and maintenance) connections for each user connection. For connections on the LIM side, the OAM connections are created on the LIM. For trunk-terminated connections, the system creates an OAM connection as well as an internal connection for a SAR exception data path on the control module. These connections on the control module use NP resources that could otherwise be used by user

connections. To maximize user connections system-wide, you can restrict the F5 OAM trunk connections the system creates. See “IP2100 memory modes” on page 1-7 for related information.

Overview of OAM-CONFIG profile settings

To restrict the number of F5 OAM trunk connections supported by the system, set the following parameter, shown with its default setting:

```
[in OAM-CONFIG]
number-of-oam-connections = 0
```

Parameter	Setting
number-of-oam-connections	Specifies the number of F5 OAM trunk connections the system creates system-wide. Valid values are from 0 through 12000. However, the system limit for OAM connections is currently 5000. With the default zero setting, the system enables OAM for all trunk-terminated connections up to a maximum of 5000 OAM connections. This results in fewer supported user connections system-wide. To allow more user connections, set this parameter to a nonzero number lower than 5000. A system reset is required to reduce the number of OAM trunk connections. However, no reset is required if you subsequently increase the number of supported OAM connections.

Required IP2100 memory mode to maximize user connections

To support the maximum number of user connections, you must set the `np-fpp-class-mem-mode` parameter to `medium-capacity`. For example:

```
admin> read system
admin> set np-fpp-class-mem-mode = medium-capacity
admin> write -f
```

For details about this setting, see “IP2100 memory modes” on page 1-7.

Additional connections created for trunk-terminated connections

For each trunk-terminated call, the system creates an additional connection for SAR exception data path. The SAR exception data path must be created for all trunk-terminated calls and cannot be limited. You should take these additional internal connections into account when provisioning trunk-terminating calls.

Notice of OAM limitation with more than 5000 user PVCs

Currently, when the number of provisioned LIM-to-trunk PVCs exceeds 5000, the maximum number of OAM connections still cannot be increased.

Configurable IP2100 port managers

The IP2100 control module supports two Gigabit Ethernet ports and a network processor (NP) capacity of approximately 3.9Gbps. An HB LIM can use up to 559Mbps (whereas other LIMs use approximately 152Mbps), which means that the total installed LIM bandwidth of the system could possibly exceed the IP2100 NP capacity. To address this possibility, the system supports user configuration of the NP port managers for each LIM slot, Gigabit Ethernet port, and trunk port. You must configure the port managers to distribute NP capacity properly as you install various modules in the system.



Note For proper operation, you must configure the NP port manager for each HB LIM in the system. For details, see “Sample HB LIM NP-PORT configuration” on page 1-13.

System-generated profiles

To enable you to configure NP port managers and to check their status, the system creates the following profiles at startup or when a port first becomes active (such as when a module is first installed):

Table 1-5. NP port manager profiles

Profile	Description
NP-PORT	Configuration profile used to request a fixed bandwidth rate. The system generates one profile for each installed LIM, one for each port on installed trunk modules (up to a maximum of two profiles per installed trunk module), and one for each Gigabit Ethernet port. For a trunk aggregation module (TRAM) the system generates two NP-PORT profiles for the total bandwidth of the TRAM. Each of the TRAM NP-PORT profiles handles a bandwidth rate equal to one OC3.
NP-PORT-STAT	Read-only status profile that indicates the actual bandwidth allocated by the NP port manager. For details about NP-PORT-STAT profiles and commands for viewing bandwidth allocation, see the <i>Stinger Administration Guide</i>.

Fixed-rate and nonfixed-rate bandwidth

Each NP port manager has a configurable fixed rate as well as a system-derived nonfixed rate, which are used together to satisfy the bandwidth requirements of traffic going out from the NP port manager on the IP2100 control module to the associated slot or port, and ultimately out to the DSL CPE (for LIMs), ATM core network (for trunk ports), or Gigabit Ethernet network (for control module GigE ports).

The aggregate bandwidth is the sum of fixed-rate and nonfixed-rate bandwidth. The fixed rate is the guaranteed bandwidth portion of aggregate bandwidth.

The NP imposes the following constraints on the selection of the fixed-rate and nonfixed-rate values for each port. Both constraints must be simultaneously met for the rate values to be valid.

- Fixed-rate + nonfixed-rate bandwidth = aggregate bandwidth

Note that the aggregate bandwidth cannot exceed the bandwidth of the interface.

- The ratio (fixed-rate/nonfixed-rate) OR (nonfixed-rate/fixed-rate) must be an integer multiple between 1 and 16.

For each port, the configured fixed rate defines the bandwidth pool that is used to satisfy the bandwidth requirements of the associated CBR, real-time VBR, and non-real-time VBR connections (in ATM) and of priority-1, priority-2, and priority-3 IP connections.

For each port, the nonfixed-rate bandwidth pool is used for UBR connections (in ATM) and of priority-0 IP connections. The system derives the amount of nonfixed-rate bandwidth from the requested fixed rate and the line capacity. The nonfixed rate is always the highest possible multiple of the requested fixed rate within the maximum line capacity.

Effect of fixed-rate configuration on oversubscription

Fixed rates are constrained by line capacities and the overall capacity of the NP. The fixed rate of a single port cannot exceed the capacity of the corresponding interface. In addition, the sum of the fixed rates allocated over all ports cannot exceed the capacity of the NP. In other words, fixed rates cannot be oversubscribed systemwide. However, oversubscription within the fixed rate bandwidth pool for a specific port is allowed—the sum of the guaranteed rates of the allocated connections can exceed the configured fixed rate by a set factor (such as 10).

Nonfixed rates are constrained only by line capacity. The nonfixed rate of a single port cannot exceed the capacity of the corresponding interface, but the sum of the nonfixed rates can exceed the overall capacity of the NP. The nonfixed rate can therefore be oversubscribed over all NP ports.

For details about connection admission control and oversubscription, see the *Stinger ATM Configuration Guide*.

Effect of fixed-rate configuration on QoS

Quality of Service (QoS) is impacted by the fixed-rate configuration in that the fixed-rate bandwidth places a limit on what is available for guaranteed bandwidth connections. For example, if you configure 100Mbps fixed-rate bandwidth on a LIM, the total guaranteed downstream QoS for all users on that LIM cannot exceed 100Mbps. You must be aware of this when configuring downstream QoS for DSL users.

Configurable NP-PORT profile settings

The NP-PORT profile contains the following parameters, shown with default values for a LIM in slot 1:

```
[in NP-PORT/{ shelf-1 slot-1 0 }]
port-address* = { shelf-1 slot-1 0 }
requested-fixed-bandwidth = 149750
fixed-bandwidth-preference = discrete-fixed-rate
```

Parameter	Setting														
port-address	Address of the LIM slot, or of the individual Gigabit Ethernet or trunk port, associated with an NP port manager. for LIM slots The address uses the following format: { shelf-M slot-N 0 } The zero port number represents all ports on the LIM. for trunk ports The address uses the following format: { shelf-1 trunk-module-M N } Each trunk port has its own configurable fixed rate. for GigE ports The address uses the following format: { shelf-1 slot-any N } The two Gigabit Ethernet ports of the IP2100 (ports 2 and 3) each have their own configurable fixed rate. The slot-any designation is used for redundancy. In the event of primary failure, the new primary control module can use the same NP-PORT profiles.														
requested-fixed-bandwidth	Requested fixed rate bandwidth in Kbps for the NP port. If allocated successfully, data will be shaped at the specified rate through the LIM, trunk, or GigE interface. Because the system must maintain a ratio between the fixed and on-fixed bandwidth on an interface, the system may not allocate all of the requested bandwidth. In addition, the setting of the fixed-bandwidth-preference parameter may limit the total interface throughput. The minimum value is 1500. The actual configurable maximum bandwidth depends on the type of the interface and the maximum NP capacity of 3.9Gbps. The maximum aggregate bandwidth (fixed and nonfixed) for each type of interface is as follows: <table> <tr> <td>HB LIM</td><td>588888Kbps</td></tr> <tr> <td>Legacy LIM</td><td>149750Kbps</td></tr> <tr> <td>OC3 TM port</td><td>149750Kbps</td></tr> <tr> <td>DS3 TM port</td><td>40641Kbps</td></tr> <tr> <td>E3 TM port</td><td>33888Kbps</td></tr> <tr> <td>OC12 TM port</td><td>599000Kbps</td></tr> <tr> <td>GigE port</td><td>1000000Kbps</td></tr> </table>	HB LIM	588888Kbps	Legacy LIM	149750Kbps	OC3 TM port	149750Kbps	DS3 TM port	40641Kbps	E3 TM port	33888Kbps	OC12 TM port	599000Kbps	GigE port	1000000Kbps
HB LIM	588888Kbps														
Legacy LIM	149750Kbps														
OC3 TM port	149750Kbps														
DS3 TM port	40641Kbps														
E3 TM port	33888Kbps														
OC12 TM port	599000Kbps														
GigE port	1000000Kbps														
fixed-bandwidth-preference	The preferred method for allocating the fixed bandwidth rate specified by the requested-fixed-bandwidth setting.														

Parameter	Setting
	discrete-fixed-rate
	This is the default setting. With this setting, the system calculates a fixed/nonfixed ratio that maximizes the aggregate throughput on the interface.
	flexible-fixed-rate
	With this setting, the system attempts to use the fixed bandwidth rate specified by the requested-fixed-bandwidth setting as closely as possible. This is the recommended setting.
	The system calculates a fixed/non-fixed ratio that maximizes the fixed-rate throughput on the interface.
	Depending on the value of requested-fixed-bandwidth , the aggregate rate with this setting may be limited to less than interface rate (line capacity).

At startup or when an interface first becomes active, the system attempts to allocate the requested bandwidth if it is available, and creates or updates the NP-PORT-STAT profile with the amount of bandwidth it was able to allocate.

Caution about modifying NP-PORT bandwidth allocations

When you configure NP-PORT profiles to request a higher fixed bandwidth for HB LIMs, you must remember the maximum network processor capacity of 3.9Gbps.

If you configure multiple bandwidth allocations which in aggregate exceed the maximum NP capacity, a bandwidth starvation condition occurs. The `info np pm bwdist` debug-level command displays fixed rate allocation, and will show bandwidth starvation conditions should they occur.



Note See Appendix A, “IP Control Module Diagnostics” for important caveats about using debug-level commands and for an overview of the `info np` command.

For example, the following command output shows bandwidth starvation on slot 12:

```
super> info np pm bwdist
NP Port Manager Rates
NP PM Bandwidth Distribution
    System Aggregate Allocated BW : 3904697
    Aggregate User Requested BW : 3977196
    Available system BW : 0
```

Port	Requested	Allocated	Status
{1 1 0}	540000	544415	BW SERVICED
{1 2 0}	9359	9393	BW SERVICED
{1 3 0}	540000	544415	BW SERVICED
{1 4 0}	9359	9393	BW SERVICED
{1 5 0}	540000	544415	BW SERVICED
{1 6 0}	9359	9393	BW SERVICED
{1 7 0}	9359	9393	BW SERVICED
{1 10 0}	9359	9393	BW SERVICED
{1 11 0}	9359	9393	BW SERVICED

{1 12 0}	540000	444733	BW STARVED
{1 13 0}	9359	9393	BW SERVICED
{1 14 0}	9359	9393	BW SERVICED
{1 15 0}	36805	36805	BW SERVICED
{1 16 0}	9359	9393	BW SERVICED
{1 17 1}	2540	2555	BW SERVICED
{1 17 2}	2540	2555	BW SERVICED
{1 18 1}	2540	2555	BW SERVICED
{1 18 2}	2540	2555	BW SERVICED
{1 0 2}	828000	833333	BW SERVICED
{1 0 3}	828000	833333	BW SERVICED

How to recover from a bandwidth starvation condition

To enable the system to recover from a bandwidth starvation condition, you must complete the following steps:

- 1 Open the NP-PORT profile for the slot or port that initiated the bandwidth starvation condition and set the `requested-fixed-bandwidth` parameter to a smaller setting.
- 2 Free up some unused bandwidth by opening the NP-PORT profile associated with an unused slot or port, and setting the `requested-fixed-bandwidth` parameter to its minimum value of 1500 (1.5Mbps).
- 3 Re-open the NP-PORT profile that initiated the bandwidth starvation condition and set the `requested-fixed-bandwidth` parameter to the setting you require.

Protection against exceeding line capacity in NP-PORT configurations

When you modify the settings of the NP-PORT profile, the system uses the new profile values to reprogram the port managers only after validating that the `requested-fixed-bandwidth` value does not exceed the maximum capacity of the line. For example, if you attempt to specify a fixed bandwidth of 200,000Kbps for a legacy LIM, the system refuses to write the NP-PORT profile and displays an error message such as the following:

```
admin> read np-port {1 1 0}
admin> set requested-fixed-bandwidth = 200000
admin> write -f
error: Port{1 1 0} fixed rate 200000 exceeds supported rate of 152576
```

If bandwidth is unavailable for allocation, the system generates log messages and the configures the port with a minimum bandwidth of 1.5Mbps for fixed rate and 16Mbps for nonfixed rate.

Sample HB LIM NP-PORT configuration

For proper operation in a capacity environment, you must explicitly configure the port bandwidth for each HB LIM. The `requested-fixed-bandwidth` value you set depends on how much traffic will be sent to the LIM. The total fixed bandwidth allocation systemwide cannot exceed the maximum NP capacity of 3.9Gbps. See “Configurable IP2100 port managers” on page 1-9 for details.

Introduction

Introduction to the Stinger IP2100 control module

For example, suppose the system supports the following modules:

```
admin> show
```

```
Controller { first-control-module } ( PRIMARY ):
```

	Reqd	Oper	Slot Type
{ shelf-1 slot-1 0 }	UP	UP	ep-72-hs-gs-ads12plus
{ shelf-1 slot-2 0 }	UP	UP	stngr-48a-ads1-card
{ shelf-1 slot-3 0 }	UP	UP	ep-72-hs-gs-ads12plus
{ shelf-1 slot-4 0 }	UP	UP	sds1-atm-v2-card
{ shelf-1 slot-5 0 }	UP	UP	ep-72-hs-gs-ads12plus
{ shelf-1 slot-6 0 }	UP	UP	sds1-atm-v2-card
{ shelf-1 slot-7 0 }	UP	UP	sds1-atm-v2-card
{ shelf-1 slot-10 0 }	UP	UP	stngr-72-shds1-card
{ shelf-1 slot-11 0 }	UP	UP	stngr-72-shds1-card
{ shelf-1 slot-12 0 }	UP	UP	ep-48-hs-gs-ads12plus
{ shelf-1 slot-13 0 }	UP	UP	ima-24t1-card
{ shelf-1 slot-14 0 }	UP	UP	ishds1-48-card
{ shelf-1 slot-15 0 }	UP	UP	ep-72-hs-gs-ads12plus
{ shelf-1 slot-16 0 }	UP	UP	ima-24t1-card
{ shelf-1 trunk-module-1 0 }	UP	UP	ds3-atm-trunk-daughter-card
{ shelf-1 trunk-module-2 0 }	UP	UP	ds3-atm-trunk-daughter-card

With this system configuration and default settings in the NP-PORT profiles, the system allocates fixed-rate bandwidth as shown in the following command output:

```
super> info np pm bwdist
```

```
NP Port Manager Rates
```

```
NP PM Bandwidth Distribution
```

```
System Aggregate Allocated BW : 432581
```

```
Aggregate User Requested BW : 433416
```

```
Available system BW : 3486580
```

Port	Requested	Allocated	Status
{1 1 0}	36805	36805	BW SERVICED
{1 2 0}	9359	9393	BW SERVICED
{1 3 0}	36805	36805	BW SERVICED
{1 4 0}	9359	9393	BW SERVICED
{1 5 0}	36805	36805	BW SERVICED
{1 6 0}	9359	9393	BW SERVICED
{1 7 0}	9359	9393	BW SERVICED
{1 10 0}	9359	9393	BW SERVICED
{1 11 0}	9359	9393	BW SERVICED
{1 12 0}	36805	36805	BW SERVICED
{1 13 0}	9359	9393	BW SERVICED
{1 14 0}	9359	9393	BW SERVICED
{1 15 0}	36805	36805	BW SERVICED
{1 16 0}	9359	9393	BW SERVICED
{1 17 1}	2540	2555	BW SERVICED
{1 17 2}	2540	2555	BW SERVICED
{1 18 1}	2540	2555	BW SERVICED
{1 18 2}	2540	2555	BW SERVICED
{1 0 2}	62500	62654	BW SERVICED
{1 0 3}	62500	62654	BW SERVICED

The following commands modify the default bandwidth setting for the HB LIMs installed in slot 1 and slot 3 to allow approximately 521Mbps of guaranteed data plus approximately 38Mbps of non-guaranteed data:

```
admin> read np-port { 1 1 }
admin> set fixed-bandwidth-preference = flexible-fixed-rate
admin> set requested-fixed-bandwidth = 521000
admin> write -f
admin> read np-port { 1 3 }
admin> set fixed-bandwidth-preference = flexible-fixed-rate
admin> set requested-fixed-bandwidth = 521000
admin> write -f
```

With the sample settings immediately above, guaranteed traffic is limited to 521Mbps on the LIM. In the absence of guaranteed traffic, the guaranteed bandwidth is used for non-guaranteed traffic, so this setting enables you to send up to 521+ 38 = 559Mbps of non-guaranteed traffic on the HB LIM.

Features not currently supported by the IP2100

The IP2100 control module does not currently support the following features:

- Control module redundancy in hosted environments
- Virtual path shaping

Network architecture overview

A Stinger IP DSLAM supports all standard Stinger ATM features, such as data and voice services over DSL. In addition, it supports IP services such as multicast video, unicast video-on-demand, and other video and IPTV applications. The services supported by the IP DSLAM are provided downstream to DSL subscribers. In the upstream direction, DSL subscribers accessing the Stinger IP DSLAM can be routed via the IP infrastructure to Internet services, or bridged/routed to a virtual LAN.

A Stinger IP DSLAM supports IP routing and QoS, ATM QoS, traffic shaping, VLAN, and multicasting capabilities to provide fast, efficient access to ATM and IP services.


Multicast video

A Stinger IP DSLAM uses the Internet Group Management Protocol (IGMP) to manage group memberships of downstream video to a PC application or set-top box, as shown in Figure 1-2. Administrators can configure levels of service that control subscribers' access to specific multicast groups. Connection to the originating router or multicast router can be across the Gigabit Ethernet interface or through a high-speed IP over ATM connection.

Introduction

Network architecture overview


Figure 1-2. Sample setup showing multicast and unicast video services



Internet and voice access

When a subscriber has DSL Integrated Access Device (IAD) equipment (such as a CellPipe®), the Stinger can deliver integrated voice and data services over the local copper loop, providing a efficient, low-cost solution for enterprise, small business, home office, and residential subscribers.


Figure 1-3. Sample setup showing Internet access and voice over ATM



Multiplexing multiple IP flows on a single ATM VCC

A Stinger IP DSLAM supports an implementation of Class of Service (CoS) that co-exists with the Stinger ATM QoS implementation. This feature allows transferring multiple IP streams (multicast and unicast) over single user-side ATM virtual circuit with different levels of priority.

Figure 1-4. Sample setup showing multiple IP flows to a CPE router



The CoS implementation enables the delivery of differentiated services over an IP infrastructure. All traffic handled by the IP control module, whether encapsulated IP or native ATM, passes through the network processor function.

Non-IP terminated ATM traffic, including operations, administration, and maintenance (OAM) F5 traffic, is treated as highest priority and handled in an *ATM pass-through mode*. This traffic passes through the network processor with no further processing.

RFC 2684 IP traffic that terminates on the IP control module is reassembled from ATM cells into IP packets. It is then classified and assigned to priority output queues. A default per-VC strict-priority queuing is supported with three priority levels as described in Table 1-6.

Table 1-6. Default classification and prioritization

Priority queue	Priority level	Packet classification assigned to queue
1	High	IP Control Protocol Classification ■ ARP/RARP protocol messages ■ ICMP protocol messages ■ RIP protocol messages ■ IGMP protocol messages
2	Medium	Multicast Classification ■ IP multicast data
3	Low	Unicast Classification ■ IP unicast data

Table 1-6 shows a default classification that occurs when no PACKET-FLOWS profile has been applied to the traffic. For information about priority queuing based on flow identification, see Chapter 5, “Ethernet and IP QoS.”

Per-VC queuing operates in conjunction with the associated ATM shaping rate. The aggregate rate of the combination of three priority queues (Class of Service Queuing with Strict Priority) associated with a particular ATM virtual circuit is controlled by the SCR (sustained cell rate) configured for the VC. In this case, SCR is configured equal to PCR (peak cell rate). Rate information is configurable in the ATM-QoS profile for each virtual circuit. For details about configuring ATM QoS, see the *Stinger ATM Configuration Guide*.

Gigabit Ethernet Configuration **2**

Configuring Gigabit Ethernet interfaces.	2-1
Configuring Gigabit Ethernet port redundancy	2-5
Configuring LACP on Gigabit Ethernet ports (IP2100 only)	2-9
Configuring STP on Gigabit Ethernet ports (IP2100 only).	2-17
Administrative tools for Gigabit Ethernet	2-25

This chapter describes how to configure the Gigabit Ethernet Media Access Control (GMAC) physical interfaces on the Stinger IP2000 and Stinger IP2100 control modules. For information about configuring the 10/100 BASE-T Ethernet management interfaces, see the *Getting Started Guide* for your Stinger platform.

For details about configuring IP interfaces on the Ethernet ports, see “Configuring IP-INTERFACE profiles for Ethernet ports” on page 4-6.

Configuring Gigabit Ethernet interfaces

The system generates an ETHERNET profile for each Ethernet port on the IP control modules.

For the IP2000 control module, the Gigabit Ethernet port has the following indexes:

```
ETHERNET/{ shelf-1 first-control-module 2 }  
ETHERNET/{ shelf-1 second-control-module 2 }
```

For the IP2100 control module, the Gigabit Ethernet ports have the following indexes:

```
ETHERNET/{ shelf-1 first-control-module 2 }  
ETHERNET/{ shelf-1 first-control-module 3 }  
ETHERNET/{ shelf-1 second-control-module 2 }  
ETHERNET/{ shelf-1 second-control-module 3 }
```

With the default settings, Gigabit Ethernet interfaces are fully operational.

Overview of ETHERNET profile settings

Following are the ETHERNET parameters, shown with default settings for a Gigabit Ethernet port:

```
[in ETHERNET/{ shelf-1 first-control-module 2 }]  
interface-address* = { shelf-1 first-control-module 2 }
```

```
link-state-enabled = no
enabled = yes
ether-group = 0
ether-if-type = utp
bridging-enabled = no
filter-name = ""
duplex-mode = full-duplex
pppoe-options = { no no "" }
bridging-options = { 0 no no transparent-bridging 0 0 "" 0 }
lACP-options = { 65535 active }
media-speed-mbit = 100mb
auto-negotiate = no
vlan-stack-tag-type = 91:00
```

For details about each of the parameters shown above, see the *Stinger Reference*.
Following are some Gigabit Ethernet-specific notes about the profile contents:

Parameter	Notes about Gigabit Ethernet settings
interface-address*	The interface-address value for a physical Gigabit Ethernet port.
link-state-enabled	With the default value, the system discards packets and does not choose an alternate route if the interface is down. If you set this to yes , the system deletes routes to the interface when the interface is unavailable, and then restores the routes when the interface becomes available again.
enabled	If you set this to no and write the profile, the interface is unavailable.
ether-group	The index of an ETHER-GROUP profile. See “ETHERNET profile settings for physical ports in a LAG” on page 2-12
bridging-enabled	See “Enabling layer 2 bridging for VLAN operations” on page 2-3.
filter-name	Name of a filter to apply to all non-VLAN traffic on the Gigabit Ethernet interface. Non-VLAN traffic consists of untagged frames and frames with VLAN IDs that do not match configured VLAN-ETHERNET profiles. See Chapter 11, “Filter Configuration.”
duplex-mode	The GMAC physical interface operates only in full-duplex mode.
pppoe-options	Not currently used by IP control modules.
bridging-options	Not currently used by IP control modules.
lACP-options	A subprofile where port-level link aggregation parameters are configured. See “Configuring LACP on Gigabit Ethernet ports (IP2100 only)” on page 2-9.
media-speed-mbit	This setting is read-only and specifies 1Gbps.

Parameter	Notes about Gigabit Ethernet settings
auto-negotiate	Enables or disables autonegotiation. Note For the IP2100 control module, you must set this value to the same value as the equipment to which it is connected.
vlan-stack-tag-type	If the layer-2 core network is using a specific EtherType value for stacked VLAN frames, you must set the <code>vlan-stack-tag-type</code> parameter to that value. For details, see “Configuring stacked VLANs” on page 3-27.

Enabling layer 2 bridging for VLAN operations

Following is the relevant parameter, shown with its default setting, for enabling bridging on the Gigabit Ethernet port:

```
[in ETHERNET/{ shelf-1 first-control-module 2 }]  
bridging-enabled = no
```

Parameter	Setting
bridging-enabled	Configures the physical port to accept Ethernet frames for bridging purposes. With the default <code>no</code> value, the system does not accept unicast packets received on this port unless the destination MAC address is equal to the MAC address of the port. Set this parameter to <code>yes</code> if the system will support VLAN operations.

The following commands enable bridging on Gigabit Ethernet:

```
admin> read ethernet { 1 8 2 }  
admin> set bridging-enabled = yes  
admin> write -f
```

Verifying the Gigabit Ethernet interface setup

To verify that the Gigabit Ethernet interface is able to handle IP traffic, first you must assign an IP address to the interface. See “Configuring IP-INTERFACE profiles for Ethernet ports” on page 4-6 for instructions.

After you assign an IP address, you can verify that the Gigabit Ethernet interface is able to handle the traffic by checking some command output. For details about the `netstat` and `gmac` commands, see the *Stinger Reference*.

You can also use the debug-level `ifmgr -d` command to verify that the Gigabit Ethernet interface is active. This is described in Appendix A, “IP Control Module Diagnostics.”

Checking the routing table

The following command output verifies that the routing table has an entry for the Gigabit Ethernet interface (IP address 100.1.1.3/32):

```
admin> netstat -rn
```

Destination	Gateway	IF	Flg	Pref	Met	Use	Age
0.0.0.0/0	1.1.2.1	ie0	SGP	60	1	3817	828
20.1.2.0/24	-	ie1-1	C	0	0	0	828
20.1.2.3/32	-	local	CP	0	0	0	828
100.0.0.0/8	-	ie1	C	0	0	4683	828
100.1.1.3/32	-	local	CP	0	0	1580	828
127.0.0.0/8	-	bh0	CP	0	0	0	828
127.0.0.1/32	-	local	CP	0	0	0	828
127.0.0.2/32	-	rj0	CP	0	0	0	828
1.1.2.0/24	-	ie0	C	0	0	1428	828
1.1.2.65/32	-	local	CP	0	0	2937	828
224.0.0.0/4	-	mcast	CP	0	0	0	828
224.0.0.1/32	-	local	CP	0	0	0	828
224.0.0.2/32	-	local	CP	0	0	0	828
224.0.0.9/32	-	local	CP	0	0	0	828
255.255.255.255/32	-	ie0	CP	0	0	0	828

Verifying the network processor setup for the interface

The network processor on the IP control module creates a connection entry for the Gigabit Ethernet interface when the interface becomes operational. You can force the network processor to create a connection entry for the Gigabit Ethernet interface by using the following command:

```
admin> gmac -n
NP setup for gmac done.
```

Verifying the SAR setup for the interface

The onboard Segmentation and Reassembly (SAR) creates an ATM connection entry for the Gigabit Ethernet interface. You can force the SAR setup by using the following command:

```
admin> gmac -s
GMAC: SAR conn. open with vpi = 0, vci = 200
```

Verifying IP packet transfer on the interface

The following command clears statistics gathered on the Gigabit Ethernet interface:

```
admin> gmac -d -c
```

The next command pings a host on the same subnet as the Gigabit Ethernet interface:

```
admin> ping 100.1.1.10
PING 100.1.1.10 (100.1.1.10): 56 data bytes
64 bytes from 100.1.1.10: icmp_seq=0 ttl=255 time=0 ms
64 bytes from 100.1.1.10: icmp_seq=1 ttl=255 time=0 ms
64 bytes from 100.1.1.10: icmp_seq=2 ttl=255 time=0 ms
64 bytes from 100.1.1.10: icmp_seq=3 ttl=255 time=0 ms
64 bytes from 100.1.1.10: icmp_seq=4 ttl=255 time=0 ms
64 bytes from 100.1.1.10: icmp_seq=5 ttl=255 time=0 ms
```

```
64 bytes from 100.1.1.10: icmp_seq=6 ttl=255 time=0 ms
64 bytes from 100.1.1.10: icmp_seq=7 ttl=255 time=0 ms
--- 100.1.1.10 ping statistics ---
8 packets transmitted, 8 packets received, 0% packet loss
round-trip min/avg/max = 0/0/0 ms
```

The following command displays GMAC statistics that show packet transfer. The txGoodPackets and rxGoodPackets fields in the command output show 8 packets transmitted and received in the ICMP sequence shown immediately above. For more details on the command output fields, see “Total transmit statistics” on page A-30 and “Total receive statistics” on page A-32.

```
admin> gmac -d
```

Gigabit Ethernet port statistics:

txOctetsLow	= 162450	rxOctetsLow	= 13192
txOctetsHigh	= 0	rxOctetsHigh	= 0
txGoodPackets	= 1874	rxGoodPackets	= 53
txPkt64	= 11	rxPkt64	= 3
txPkt65127	= 1813	rxPkt65127	= 0
txPkt128255	= 0	rx128255	= 0
txPkt256511	= 50	rx256511	= 50
txPkt5121023	= 0	rx5121023	= 0
txPkt1024Max	= 0	rx1024Max	= 0
txPktDefer	= 0	rxMacType	= 0
txPktUndSz	= 0	rxCrcErrors	= 0
txUnderFlow	= 0	rxUnderSize	= 0
txPfcf	= 0	rxOverSize	= 0
txPfcc	= 0	rxAlmostFull	= 0
txRfcf	= 0	rxOverRun	= 0
txRfcc	= 0	rxMulticastPkts	= 0
txOverflow	= 0	rxBroadcastPkts	= 50
txAlmostFull	= 0	rxJabber	= 0
		rxPfc	= 0
		rxRfc	= 0

Configuring Gigabit Ethernet port redundancy

Systems with redundant IP control modules can be configured to enable Gigabit Ethernet port redundancy. With proper configuration, RFC 2684 (MPoA) connections, VLAN connections, and LAN MBONE interface functions can be maintained across the Gigabit Ethernet interface in the event of primary control module failure.



Note A soft IP interface configuration is required for Gigabit Ethernet redundancy of RFC 2684 connections and a LAN MBONE interface.

Configuring a soft IP interface for Gigabit Ethernet redundancy

The soft IP interface is an internal interface that is not associated with a specific physical port, but that can be accessed through the Ethernet interface of whichever controller is primary. For background information, see “Defining a soft interface for increased accessibility” on page 4-8.

You can use the default (system-created) soft IP interface for Gigabit Ethernet redundancy, which uses the following index:

IP-INTERFACE/{ { 0 0 0 } 0 }

However, if you have already used the default profile for the soft IP address of the 10/100M base Ethernet management ports, you can create another soft IP interface using a profile index of { { 0 0 0 } x }, as long as the IP address in that profile is on the same subnet as the Gigabit Ethernet ports.




Note The system associates its Ethernet interfaces with a particular soft address based on the subnet assignment. The IP interface address of the Gigabit Ethernet ports on the primary and secondary controllers and the soft IP interface address must be on the same subnet.

Configuring Gigabit Ethernet redundancy for RFC 2684 (IPoA) connections

Figure 2-1 shows a Stinger with redundant IP control modules. The Gigabit Ethernet port in slot 8 ({ { 1 8 2 } 0 }), the Gigabit Ethernet port in slot 9 ({ { 1 9 2 } 0 }), and the soft IP interface ({ { 0 0 0 } 1 }), all have IP address assignments on the same subnet. In addition, the external router has a routing table entry that specifies the soft IP interface address as the gateway to the CPE router destination.

Figure 2-1. Gigabit Ethernet redundancy for RFC 2684 connectivity



The following commands configure the Gigabit Ethernet port in slot 8:

```
admin> read ip-interface { { 1 8 2 } 0 }
admin> set ip-address = 10.99.99.101/24
admin> write -f
```

The next commands configure the Gigabit Ethernet port on slot 9:

```
admin> read ip-interface { { 1 9 2 } 0 }
admin> set ip-address = 10.99.99.201/24
admin> write -f
```

The following commands configure a soft IP interface on the same subnet:

```
admin> new ip-interface { { 0 0 0 } 1 }
admin> set ip-address = 10.99.99.100/24
admin> write -f
```

When you write the profile of the soft interface, the system displays a message:

```
LOG notice, Shelf 1, Controller-1, Time: 11:42:57--
Soft ip will be effective if the ip-addr of primary controller is
configured.
```

To ensure that the external router can reach the CPE router in Figure 2-1, the external router must specify the soft IP address as the gateway to the CPE router destination address. For example,

```
destination-address = 10.14.16.1/24
gateway-address = 10.99.99.100/24
```

Configuring Gigabit Ethernet redundancy for VLAN bridging

To enable the system to maintain VLAN bridging connections in the event of the primary control module failing, you must configure the VLAN on the soft interface, using the expression `any-slot` or `0` as the slot number. For example, the following commands define a new Gigabit Ethernet-redundant VLAN with bridge group 95 and VLAN ID 95:

```
admin> new bridge-group 95
admin> set enable = yes
admin> set bridging-group = 95
admin> set mac-entry-age-time = 300
admin> set igmp-snooping-enabled = yes
admin> set port-block-enabled = yes
admin> set lan-router-interface-address physical-address shelf = shelf-1
admin> set lan-router-interface-address physical-address slot = any-slot
admin> set lan-router-interface-address physical-address item-number = 2
admin> set lan-router-interface-address logical-item = 95
admin> write -f
admin> new vlan-ethernet { { 1 a 2 } 95 }
admin> set enabled = yes
admin> set bridging-options bridging-group = 95
admin> set bridging-options bridge = yes
admin> set bridging-options bridge-type = transparent-bridging
admin> write -f
```

The following commands define a CONNECTION profile that uses the redundant interface, so it will be maintained across a switchover:

```
admin> new connection raj-agg-1
admin> set active = yes
admin> set encapsulation-protocol = atm
```

```
admin> set ip-options ip-routing-enabled = no
admin> set bridging-options bridging-group = 95
admin> set bridging-options bridge = yes
admin> set bridging-options bridge-type = transparent-bridging
admin> set atm-options vci = 95
admin> set atm-options nailed-group = 151
admin> write -f
```

To modify an existing VLAN for Gigabit Ethernet redundancy, you must create a new configuration and then delete the old one. For example, the following command shows an existing VLAN with ID 50:

```
admin> dir vlan
37 07/21/2005 17:38:24 { { shelf-1 first-control-module 2 } 50 }
```

The next commands modify the VLAN to enable Gigabit Ethernet redundancy for the connection:

```
admin> read vlan { { 1 8 2 } 50 }
admin> set interface-address = { { 1 0 2 } 50 }
(New index value; will save as new profile VLAN-ETHERNET/{ { shelf-1 any-
slot 2 } 50 }.)
admin> write -f
```

The following commands list and then delete the older VLAN-ETHERNET profile:


```
admin> dir vlan
37 08/15/2005 09:00:30 { { shelf-1 any-slot 2 } 50 }
37 07/21/2005 17:38:24 { { shelf-1 first-control-module 2 } 50 }
admin> delete vlan { { 1 8 2 } } 50}
```

Configuring a redundant LAN MBONE

To support redundancy for a LAN MBONE interface, you must configure a soft IP interface for the Gigabit Ethernet ports, enable multicast on both ports, and use the any-slot expression in the `mbone-lan-interface` parameter setting.

Figure 2-2 shows a Stinger with redundant IP control modules. The Gigabit Ethernet port in slot 8 ({ { 1 8 2 } 0 }), the Gigabit Ethernet port in slot 9 ({ { 1 9 2 } 0 }), and the soft IP interface ({ { 0 0 0 } 1 }), all have IP address assignments on the same subnet and both physical ports enable multicast.

Figure 2-2. Gigabit Ethernet redundancy for a LAN MBONE



The following commands configure the Gigabit Ethernet port in slot 8:

```
admin> read ip-interface { { 1 8 2 } 0 }
admin> set ip-address = 10.99.99.101/24
admin> set multicast-allowed = yes
admin> write -f
```

The next commands configure the Gigabit Ethernet port on slot 9:

```
admin> read ip-interface { { 1 9 2 } 0 }
admin> set ip-address = 10.99.99.201/24
admin> set multicast-allowed = yes
admin> write -f
```

The following commands configure a soft IP interface on the same subnet:

```
admin> new ip-interface { { 0 0 0 } 1 }
admin> set ip-address = 10.99.99.100/24
admin> write -f
```

The following commands enable the multicast forwarding function and specify a redundant LAN MBONE configuration:

```
admin> read ip-global
admin> set multicast-forwarding = yes
admin> set multiple-mbone mbone-lan-interface 1 = { { 1 0 2 } 0 }
admin> write -f
admin> list multiple-mbone mbone-lan-interface
[in IP-GLOBAL:multiple-mbone:mbone-lan-interface]
mbone-lan-interface[1] = { { shelf-1 any-slot 2 } 0 }
mbone-lan-interface[2] = { { any-shelf any-slot 0 } 0 }
mbone-lan-interface[3] = { { any-shelf any-slot 0 } 0 }
mbone-lan-interface[4] = { { any-shelf any-slot 0 } 0 }
```



Note With this configuration, the LAN MBONE is supported on the Gigabit Ethernet port of the controller in slot 8 or slot 9, whichever is primary. Following a switchover, each IGMP client must rejoin its group to receive multicast traffic.

Configuring LACP on Gigabit Ethernet ports (IP2100 only)

You can bundle the bandwidth of the IP2100 dual Gigabit Ethernet ports into one logical link by configuring IEEE 802.3ad Link Aggregation Control Protocol (LACP). Link aggregation provides the following benefits:

- **Increased bandwidth**
The capacity of two GigE links is combined into one logical link. Both links must operate at the same data rate in full-duplex mode.
- **Increased availability**
The failure or replacement of a single link within a Link Aggregation Group (LAG) does not affect service from a client perspective.
- **Load sharing**
Traffic is distributed across both links.

LACP configuration overview

To configure the IP2100 Gigabit Ethernet ports to use link aggregation, you must complete the following steps:

- 1 Create an ETHER-GROUP profile for the LAG. The profile is indexed by a unique number. For example:

```
ETHER-GROUP/1
```

- 2 Configure a virtual ETHERNET profile for the LAG. The profile index is a virtual physical address that includes the special slot designation `slot-ether-group` and the index number of the ETHER-GROUP profile. For example, the index of a virtual ETHERNET profile for ETHER-GROUP 1 is as follows:

```
ETHERNET/{ shelf-1 slot-ether-group 1 }
```

- 3 Specify the ETHER-GROUP index, and if necessary configure LACP options, in the ETHERNET port profiles.

```
ETHERNET/{ shelf-1 first-control-module 2 }
```

```
ETHERNET/{ shelf-1 first-control-module 3 }
```

- 4 If required, provide system-level LACP information in the following profile:

```
LACP
```

ETHER-GROUP settings

Following are the ETHER-GROUP profile contents for link-aggregated ports, shown with default settings:

```
[in ETHER-GROUP/1]
index-item-number* = 1
enabled = yes
type = lacp
```

Parameter	Setting
<code>index-item-number</code>	The number of the group. You must assign a nonzero value to this field.
<code>enabled</code>	Enables or disables the profile. The profile is enabled by default. You can disable a profile to disable link aggregation of the individual links in a LAG.
<code>type</code>	The type of group. The default <code>lacp</code> type is required for configuring 802.3ad LACP.

Virtual ETHERNET profile settings for a LAG

The virtual ETHERNET profile representing a LAG is an operator-created ETHERNET profile indexed by a special virtual physical address.

When the Gigabit Ethernet ports are operating as a LAG, all port-specific settings are taken from the virtual ETHERNET profile rather than the ETHERNET profiles for the physical ports. This ensures that the two ports will have identical configurations, and that any changes to the configuration can be made in one place only.

When you create a virtual ETHERNET profile, some default values are not applicable for Gigabit Ethernet ports. Following are the default settings in a new ETHERNET profile:

```
[in ETHERNET/{ any-shelf any-slot 0 } (new)]
interface-address* = { any-shelf any-slot 0 }
link-state-enabled = no
enabled = yes
ether-group = 0
ether-if-type = utp
bridging-enabled = no
filter-name = ""
duplex-mode = full-duplex
pppoe-options = { no no "" }
bridging-options = { 0 no no transparent-bridging 0 0 "" 0 }
lacp-options = { 65535 active }
media-speed-mbit = 100mb
auto-negotiate = no
vlan-stack-tag-type = 91:00
```

Parameter	Setting
interface-address	For an ETHERNET profile representing a LAG, you must specify a virtual physical address ({ shelf slot item }) with the following settings: shelf The shelf number. slot slot-ether-group item The index of an ETHER-GROUP profile. For example, the address of a LAG ETHERNET profile associated with ETHER-GROUP 1 is: { shelf-1 slot-ether-group 1 }
link-state-enabled	A setting (not required) for deleting routes to a link when the link state is down.
enabled	Enables or disables the profile.
ether-group	Not used in a virtual ETHERNET profile representing a LAG.
ether-if-type	Not used in a virtual ETHERNET profile representing a LAG.
bridging-enabled	Enables or disables layer 2 bridging for VLAN operations on the LAG. See “Overview of ETHERNET profile settings” on page 2-1.
filter-name	Name of a filter to apply to all non-VLAN traffic on the Gigabit Ethernet interface. See “Overview of ETHERNET profile settings” on page 2-1.
duplex-mode	For a virtual ETHERNET profile representing a LAG, this value must be set to full-duplex .
pppoe-options	Not currently used by IP control modules.

Gigabit Ethernet Configuration

Configuring LACP on Gigabit Ethernet ports (IP2100 only)

Parameter	Setting
bridging-options	Not currently used by IP control modules.
lacp-options	Not used in a virtual ETHERNET profile representing a LAG.
media-speed-mbit	For a virtual ETHERNET profile representing a LAG, this value must be set to 1000mb .
auto-negotiate	For the IP2100 control module, you must set this value to the same value as the equipment to which it is connected.
vlan-stack-tag-type	If the layer-2 core network is using a specific EtherType value for stacked VLAN frames, you must set the vlan-stack-tag-type parameter to that value. See “Overview of ETHERNET profile settings” on page 2-1.

ETHERNET profile settings for physical ports in a LAG

To attach two Gigabit Ethernet ports to a LAG, the following settings are required in each port's ETHERNET profile. When link aggregation is in use, these parameters (shown with default settings) are the only settings taken from these profiles:

```
[in ETHERNET/{ shelf-1 first-control-module 2 }]
interface-address* = { shelf-1 first-control-module 2 }
enabled = yes
ether-group = 0
lacp-options = { 65535 active }

[in ETHERNET/{ shelf-1 first-control-module 2 }:lacp-options]
priority = 65535
participation = active
```

Every other parameter from the port ETHERNET profile is ignored and the values found in the LAG ETHERNET profile are used instead.

Parameter	Setting
enabled	Enables or disables the port. You can set this parameter to no to bring the port down without detaching it from the LAG.

Parameter	Setting				
ether-group	The index of an ETHER-GROUP profile. When set to 0 (default), the port operates individually. If set to a nonzero value, the system searches for an ETHER-GROUP profile indexed by this value. If it finds one, the port will be attached to that group provided that the following conditions are true: ■ The ETHER-GROUP profile is enabled. ■ The ETHER-GROUP profile specifies a type of <code>lacp</code> (the only type currently supported). ■ A valid ETHERNET profile has been created for the LAG. If the port cannot be attached it will continue to operate as an individual port. You can use this setting in the port profile to detach a port from a LAG or attach it to a new one.				
lacp-options	A subprofile where port-level link aggregation parameters are configured. There are currently two LACP parameters: <table> <tr> <td>priority</td><td>802.3ad port priority. The higher the number, the lower the priority. A configured port priority may be useful to the remote LACP peer for arbitration between aggregation candidates if the peer has a restriction on the number of ports in a LAG.</td></tr> <tr> <td>participation</td><td>802.3ad participation (active or passive). With the default <code>active</code> value, the port will initiate LACP negotiation. If set to <code>passive</code>, the port only responds to LACP negotiation.</td></tr> </table>	priority	802.3ad port priority. The higher the number, the lower the priority. A configured port priority may be useful to the remote LACP peer for arbitration between aggregation candidates if the peer has a restriction on the number of ports in a LAG.	participation	802.3ad participation (active or passive). With the default <code>active</code> value, the port will initiate LACP negotiation. If set to <code>passive</code> , the port only responds to LACP negotiation.
priority	802.3ad port priority. The higher the number, the lower the priority. A configured port priority may be useful to the remote LACP peer for arbitration between aggregation candidates if the peer has a restriction on the number of ports in a LAG.				
participation	802.3ad participation (active or passive). With the default <code>active</code> value, the port will initiate LACP negotiation. If set to <code>passive</code> , the port only responds to LACP negotiation.				

LACP profile settings

The system LACP profile has only one parameter, shown with its default value;
[in LACP]
system-priority = 65535

Parameter	Setting
system-priority	802.3ad system priority. The higher the number, the lower the priority. A configured system priority may be useful to the remote LACP peer for arbitration between aggregation candidates if the peer has a restriction on the number of ports in a LAG.

Sample link aggregation configuration

The following sample configuration bundles the capacity of the two IP2100 Gigabit Ethernet ports into a 2Gb LAG.

- 1 Create an ETHER-GROUP profile for the LAG. For example:

```
admin> new ether-group
admin> set index-item-number = 1
admin> set enabled = yes
admin> set type = lacp
admin> write -f
```
- 2 Create a new virtual ETHERNET profile for the LAG. For example:

```
admin> new ethernet { 1 slot-ether-group 1 }
admin> set ether-group = 1
admin> set media-speed-mbit = 1000mb
admin> set auto-negotiate = yes
admin> write -f
```
- 3 Modify the system-generated ETHERNET profiles for the two ports. For example:

```
admin> read ethernet { 1 first-control-module 2 }
admin> set enabled = yes
admin> set ether-group = 1
admin> write -f
admin> read ethernet { 1 first-control-module 3 }
admin> set enabled = yes
admin> set ether-group = 1
admin> write -f
```

Configuring a routed VLAN on the aggregated bandwidth

The following sample configuration creates a routed VLAN that uses the 2Gb LAG.

- 1 Create a VLAN-ETHERNET profile. Specify the index of the virtual ETHERNET profile as the physical address. For example:

```
admin> new vlan-ethernet
admin> set interface-address physical-address shelf = shelf-1
admin> set interface-address physical-address slot = slot-ether-group
admin> set interface-address physical-address item-number = 1
admin> set interface-address logical-item = 555
admin> set enabled = yes
admin> set bridging-options bridge-type = no-bridging
admin> write -f
```
- 2 Create a virtual IP interface for the routed VLAN. Specify the index of the virtual ETHERNET profile as the physical address. For example:

```
admin> new ip-interface
```

```
admin> set interface-address physical-address shelf = shelf-1
admin> set interface-address physical-address slot = slot-ether-group
admin> set interface-address physical-address item-number = 1
admin> set interface-address logical-item = 555
admin> set ip-address = 185.1.1.10/24
admin> set multicast-allowed = yes
admin> set vlan-enabled = yes
admin> set vlan-id = 555
admin> write -f
```

Configuring an N:1 bridged VLAN that uses the aggregated bandwidth

The following sample configuration creates a routed VLAN that uses the 2Gb LAG.

- 1 Create a VLAN-ETHERNET profile. Specify the index of the virtual ETHERNET profile as the physical address. For example:

```
admin> new vlan-ethernet
admin> set interface-address physical-address shelf = shelf-1
admin> set interface-address physical-address slot = slot-ether-group
admin> set interface-address physical-address item-number = 1
admin> set interface-address logical-item = 1000
admin> set enabled = yes
admin> set bridging-options bridging-group = 1000
admin> set bridging-options bridge = yes
admin> set bridging-options bridge-type = transparent-bridging
admin> write -f
```
- 2 Create a bridge group for the VLAN. Specify the index of the virtual ETHERNET profile as the physical address portion of the “router” interface address. For example:

```
admin> new bridge-group
admin> set enable = yes
admin> set bridging-group = 1000
admin> set igmp-snooping-enabled = yes
admin> set lan-router-interface physical-address shelf = shelf-1
admin> set lan-router-interface physical-address slot = slot-ether-group
admin> set lan-router-interface physical-address item-number = 1
admin> set lan-router-interface logical-item = 1000
```

Configuring an MBONE interface on the aggregated ports

The following commands configure an MBONE interface that uses the aggregated GigE bandwidth and the routed VLAN interface created in “Configuring a routed VLAN on the aggregated bandwidth” on page 2-14:

```
admin> read ip-global
```

```
admin> set multicast-forwarding = yes
admin> set multiple-mbone mbone-lan 1 physical shelf = shelf-1
admin> set multiple-mbone mbone-lan 1 physical slot = slot-ether-group
admin> set multiple-mbone mbone-lan 1 physical item-number = 1
admin> set multiple-mbone mbone-lan 1 logical-item = 555
admin> write -f
```

For details about MBONE configuration, see “Configuring MBONE interfaces” on page 9-3.

LACP implementation details

The information in this section relates to the IEEE 802.3ad specification, and provides details that may be useful in communicating with a remote LACP peer.

- **Distribution and collection of frames**

The Stinger LACP implementation has coupled control, which means that the system does not separate distributing and collecting functions.

- **Configuration capabilities and restrictions**

Only the two Gigabit Ethernet ports on an IP2100 control module support LACP and can be part of a LAG. Beyond this physical limitation, there is no limitation on the number of ports in a LAG. For example, some implementations might support LACP on ten ports but restrict the number of ports in a LAG to eight. Because the system does not enforce any such limitation, the arbitration mechanisms from 802.3 section 43.6 do not apply.

- **Admin and Operational keys**

The Stinger LACP implementation is static in the sense that Gigabit Ethernet ports are not automatically and dynamically aggregated. LACP only runs on ports that have been configured to be part of a LAG.

Ports that are part of a LAG have the same Admin Key by construction, but they may not have the same Operational Key. For example, one of the ports might have a cable disconnect or a cabling error. In this case, only the first port to be aggregated will be part of the configured LAG, and the other port will be unusable until the operational error is corrected.

- **Addressing**

Because the Stinger LACP implementation is static, the LAG MAC address does not fluctuate as ports go up and down in the LAG. The LAG takes the MAC address of its lowest numbered port. You can view the MAC address assigned to the LAG in the LAG ETHER-INFO profile.

- **Distribution algorithm**

The Stinger LACP implementation uses the IP header checksum algorithm described in RFC 1071, *Computing the Internet Checksum*. Of course, the checksum is only relevant provided that both ports are distributing. If one of the two links is down, all traffic uses the other link.

- For bridged traffic, the checksum is calculated based on the MAC source address plus the MAC destination address.
- For routed traffic, the checksum is calculated based on the IP source address plus the IP destination address.

- Limitations
 - Currently, the Link Aggregation Marker Protocol (LAMP) is not supported.
 - Single-port LAG is not supported. (You cannot create a LAG that contains a single port.)
 - LACP cannot be used in conjunction with L2TP.

LACP-related diagnostics

The following commands provide information about link aggregation on a GE-OLIM:

- The `lacp` command displays information about control traffic and traffic that terminates on or originates from the Stinger host. It includes a `wizard` option that diagnoses LACP misconfigurations. For details, see the *Stinger Reference*.
- The `stats` command supports a new `lacp` option for displaying basic LACP distribution statistics for upstream traffic, and for clearing LACP statistics counters. For details, see the *Stinger Reference*.
- The debug-level `info np` command now displays some LACP-related information. For details, see “Network processor-related diagnostics” on page A-20.

Configuring STP on Gigabit Ethernet ports (IP2100 only)

IP2100 systems support the Rapid Spanning Tree Protocol (RSTP), which is a Layer 2 link management protocol that prevents loops in a switched network with redundant data paths.


The protocol creates a spanning tree using a subset of the links, but ensuring that all switches are reachable within the tree. Spanning tree ensures that there is only one path between any pair of switches. This avoids the receipt of duplicate messages by the switches, which could otherwise occur if a loop exists in the network. Frames are forwarded only on links that make up the spanning tree, thereby ensuring that frames are not forwarded in endless loops. If there is a change in the topology (for example, when switches or links become available or unavailable), the protocol reconfigures the tree.

RSTP (IEEE 802.1w) is an evolution of the Spanning Tree Protocol (STP) defined in the IEEE 802.1d standard. With this software version, only one RSTP instance is supported. That is, only one spanning tree is built. Future versions will allow multiple instances of STP with the Multiple Spanning Tree Protocol (MSTP). For consistency with future versions, the descriptions in this document refer to RSTP simply as *STP*.

Primary application with this software version

With this software version, the primary application of STP on IP2100 systems is to enable redundant uplinks to the IP network for VLAN configurations. In Figure 2-3, the two STP-enabled Gigabit Ethernet (GigE) ports of the IP2100 control module are linked to two Ethernet switches to provide redundant uplinks.

Figure 2-3. Redundant uplinks for a VLAN



Limitations with this software version

1:1 VLAN circuit connections and stacked VLAN connections are not currently supported on STP-enabled ports. Only the following connection types currently work on STP-enabled ports:

- Routed connections
- Transparent bridged connections

STP recognizes ports but does not recognize VLANs. For this reason, the following restrictions apply:

- To provide fully redundant uplinks to VLANs, both IP2100 GigE ports must support the same VLAN configurations.

When you create a VLAN-ETHERNET profile with the special `slot-ether-group` designation in the index, the system automatically creates the specified VLAN interface for both (grouped) GigE ports. The STP subsystem decides which port to use for forwarding upstream traffic.

- Load-balancing is not supported at the VLAN level.

Currently, you cannot have one of the GigE ports in a forwarding state for one set of VLANs and in a discarding state for another set of VLANs.

- Subtending by switching between the GigE ports is supported only for transparent bridged connections.

Configuration overview

STP configuration requires both bridge-level settings for the grouped IP2100 GigE ports, and port-level settings for each of the grouped ports. To configure STP, you must complete the following steps:

- 1 Group the set of IP2100 GigE ports that will participate in STP.
You group the ports by creating an ETHER-GROUP profile and specifying the index of that profile in each port's ETHERNET profile.
Note that with this software version, only one ETHER-GROUP profile is supported for STP, because only one spanning tree instance is built.
- 2 Configure bridge-level STP options for this group.
The bridge-level parameters are specified in the ETHER-GROUP profile.
- 3 Configure port-level STP options for each individual port.
The port-level parameters are specified in each port's ETHERNET profile.
- 4 Configure VLANs on the grouped ports.

When you create a VLAN-ETHERNET profile with an index in the following format:

```
VLAN-ETHERNET { { shelf-1 slot-ether-group ether-group-ID } VLAN-ID }
```

The system creates a VLAN interface with the specified VLAN-ID on both STP-enabled ports. Instead of creating up to 4096 VLAN interfaces, the system can now create twice that number, because for each VLAN interface only one of the ports will be active. VLAN IDs must be unique among the STP-enabled ports.

When the configuration is complete, the STP subsystem determines whether one or both of the IP2100 GigE ports will forward data traffic, and begins interacting with other STP bridges by transmitting and receiving STP control frames (BPDUs).

Bridge-level STP settings in the ETHER-GROUP profile

Following are the ETHER-GROUP profile contents, shown with default settings, for configuring bridge-level STP settings:

```
[in ETHER-GROUP/1]
index-item-number* = 1
enabled = yes
type = lacp
STP-options = { stp 32768 20 2 15 }
[in ETHER-GROUP/1:STP-options]
protocol = stp
bridge-priority = 32768
max-age = 20
hello-time = 2
forward-delay = 15
```

Parameter	Setting
index-item-number	The number of the group. You must assign a nonzero value to this field. With this software version, only one ETHER-GROUP profile is supported for STP.
enabled	Enables or disables the profile. The profile is enabled by default. You can disable a profile to disable STP.
type	The type of group. Set this parameter to <code>stp</code> .
protocol	The type of Spanning Tree Protocol. Currently, <code>stp</code> is the only supported setting.

Parameter	Setting
bridge-priority	The 802.1D Bridge Priority value. Bridge priority forms part of the bridge identifier advertised to other bridges. It is used to determine the root bridge. The lower the number, the higher the priority. The default setting is 32768. The valid range is from 0 to 61,440. For STP, the priority value must be a multiple of 4096. This allows the setting of 16 relative priority values. You can use this setting to control the designation of the system as root bridge. Although it is not recommended to use the Stinger system as root bridge because doing so might affect the QoS of existing connections, the system does have the capability to act as root bridge.
max-age	The 802.1D Bridge Max Age value. This is the maximum age (in seconds) of received protocol information before it is discarded. The default is 20 seconds. The valid range is from 6 to 40.
hello-time	The 802.1D Bridge Hello Time value. This is the interval (in seconds) between transmission of periodic Configuration BPDUs by designated ports. Valid settings are 1 and 2. The default is two seconds.
forward-delay	The 802.1D Bridge Forward Delay value. This is the time (in seconds) spent by a port in Listening State and Learning State before moving to Learning or Forwarding State respectively. The default value is 15 seconds. The valid range is from 4 to 30.

Port-level STP settings in each port's ETHERNET profile

Following are the parameters, shown with default settings for port 2, for configuring port-level STP settings on the IP2100 GigE ports:

```
[in ETHERNET/{ shelf-1 first-control-module 2 }]
enabled = yes
ether-group = 0
stp-options = { 128 0 no point-to-point }
[in ETHERNET/{ shelf-1 first-control-module 2 }:stp-options]
port-priority= 128
path-cost = 0
edge-port = no
link-type = point-to-point
```

Parameter	Setting
enabled	Enables or disables the port for use.

Parameter	Setting
ether-group	The index of an ETHER-GROUP profile. When set to 0 (default), the port operates individually. If set to a nonzero value, the system searches for an ETHER-GROUP profile indexed by this value. If it finds one, the port will be attached to that group provided that the following conditions are true: ■ The ETHER-GROUP profile is enabled. ■ The ETHER-GROUP profile specifies a type of stp. ■ The ETHERNET profile is STP-enabled. If the port cannot be attached it will continue to operate as an individual port.
port-priority	A priority value, from 1 to 255 in multiples of 16, used to give preference to one port over another. The lower the value, the higher the priority. The default setting is 128. The valid range is 1-255, but for STP the value must be a multiple of 16.
path-cost	When two interfaces on a switch are part of a loop, the spanning-tree port priority and path cost settings determine which interface is put in the forwarding state and which is put in the blocking state. The path cost of the port, which is used in computing the root path cost (the cost from the port to the root switch). The path cost is related to link speed. With the default zero value, the system itself chooses the default value based on the link speed. Defaults for different link speeds are listed in Table 17.3 of the IEEE standard 802.1D-2004.
edge-port	When two interfaces on a switch are part of a loop, the spanning-tree port priority and path cost settings determine which interface is put in the forwarding state and which is put in the blocking state. Whether the port is at the edge of the switched network (true or false). This parameter should be set to true only if the port is connected to host machines, and not to other switches. If the port is connected to one or more other switches, use the default false setting.
link-type	Type of link between the port and the next switch or host machine. Valid values are point-to-point and shared.

Sample configuration with transparent bridging

To enable a bridge group to operate across the STP-enabled ports, the “router” interface for the bridge group must be supported on both ports. To allow this, use the `slot-ether-group` designation as the LAN router slot.

This sample configuration enables the IP2100 GigE ports to participate in STP. It also shows the use of the STP command, which provides detailed information about the STP subsystem. For information about this command, see the *Stinger Reference*.

- 1 Create an ETHER-GROUP profile and enable STP. For example:

```
admin> new ether-group
admin> set index-item-number = 2
admin> set type = stp
admin> write -f
```

- 2 Display the bridge-level settings to verify that STP is enabled. For example:

```
admin> stp -b
Bridge Level STP information:
-----
STP Enabled: Yes
Protocol: Stp
Ether-Group: 2
Bridge Priority: 32768
Max Age: 20
Hello Time: 2
Forward Delay: 15
```

This command shows that STP is enabled and a non-zero ETHER-GROUP number has been assigned. The other fields are bridge-level STP configuration parameters, which use the default values in this example. See the *Stinger Reference* for more detail.

- 3 Display the STP state machine. For example:

```
admin> stp -z
% TAOS: Spanning Tree Enabled - Learning Enabled
% TAOS: Ageing Time 300 - Root Path Cost 0 - Priority 32768
% TAOS: Forward Delay 15 - Hello Time 2 - Max Age 20 - Root Port 0
% TAOS: Root Id 8000000000000000
% TAOS: Bridge Id 8000000000000000
% TAOS: 0 topology changes - last topology change 0%
```

Note that the state machines has information regarding only one bridge, which is named TAOS. See the *Stinger Reference* for more detail.

- 4 Open the ETHERNET profile for the first GigE port and specify the index of the ETHER-GROUP profile. This enables the port to participate in STP. For example:

```
admin> read ethernet { 1 8 2 }
admin> set ether-group = 2
admin> write -f
```

- 5 Repeat step 4 for the second GigE port. For example:

```
admin> read ethernet { 1 8 3 }
admin> set ether-group = 2
```

```
admin> write -f
```

- 6 Display the port-level settings. For example:

```
admin> stp -p
PhyAddr  Prio  State      Cost    edge  LinkType  initState
{1 8 2}  128    Forwarding 20000   No    Point-to-Point Running
{1 8 3}  128    Forwarding 20000   No    Point-to-Point Running
```

- 7 Configure a VLAN-ETHERNET profile using the group index.

This command shows that the two GigE ports are running STP, because the `initState` column displays `Running`. If this column displays `Configured`, it indicates that the port is configured for STP but that STP is not running, perhaps because the link state of the port is down, or because the system did not find an interface for the port. See the *Stinger Reference* for more detail.

- 8 At this point, display the STP state machine again. Now the command shows that state machine information for both the ports `ie1` and `ie2` participating in STP. See the *Stinger Reference* for more detail. For example:

```
admin> stp -z
% TAOS: Spanning Tree Enabled - Learning Enabled
% TAOS: Ageing Time 300 - Root Path Cost 0 - Priority 32768
% TAOS: Forward Delay 15 - Hello Time 2 - Max Age 20 - Root Port 0
% TAOS: Root Id 800000d052025f56
% TAOS: Bridge Id 800000d052025f56
% TAOS: 2 topology changes - last topology change 6899%
% ie2: Port 2 - Id 8002 - Role Designated - State Forwarding
% ie2: Configured path cost 20000 - Designated path cost 0
% ie2: Designated port id 8002 - Priority 128
% ie2: Designated Root 800000d052025f56
% ie2: Designated Bridge 800000d052025f56
% ie2: Message Age 0 - Max Age 20
% ie2: Hello Time 2 - Forward Delay 15
% ie2: Forward Timer 0 - Msg Age Timer 0 - Hello Timer 2
% ie2: Version Rapid Spanning Tree Protocol - Received None - Send STP
% ie2: No portfast configured - Current portfast off
% ie2: Configured Link Type point-to-point - Current point-to-point
% ie2: forward-transitions 0
%
% ie1: Port 1- Id 8001 - Role Designated - State Forwarding
% ie1: Configured path cost 20000 - Designated path cost 0
% ie1: Designated port id 8001 - Priority 128
% ie1: Designated Root 800000d052025f56
% ie1: Designated Bridge 800000d052025f56
% ie1: Message Age 0 - Max Age 20
% ie1: Hello Time 2 - Forward Delay 15
% ie1: Forward Timer 0 - Msg Age Timer 0 - Hello Timer 1
% ie1: Version Rapid Spanning Tree Protocol - Received None - Send STP
% ie1: No portfast configured - Current portfast off
% ie1: Configured Link Type point-to-point - Current point-to-point
% ie1: forward-transitions 0
```

- 9 Create a VLAN on the STP-enabled ports.

For example, the following commands create VLAN 10, which belongs to BRIDGE-GROUP 10. Note that `bridge-type` is set to `transparent-bridging`.

Gigabit Ethernet Configuration

Configuring STP on Gigabit Ethernet ports (IP2100 only)

```
admin> new bridge-group
admin> set enable = yes
admin> set bridging-group = 10
admin> set igmp-snooping-enabled = yes
admin> set lan-router-interface physical-address shelf = shelf-1
admin> set lan-router-interface physical-address slot = slot-ether-group
admin> set lan-router-interface physical-address item-number = 1
admin> set lan-router-interface logical-item = 10
admin> write -f
admin> new vlan-ethernet
admin> set interface-address physical-address shelf = shelf-1
admin> set interface-address physical-address slot = slot-ether-group
admin> set interface-address physical-address item-number = 1
admin> set interface-address logical-item = 10
admin> set enabled = yes
admin> set bridging-options bridging-group = 10
admin> set bridging-options bridge = yes
admin> set bridging-options bridge-type = transparent-bridging
admin> write -f
```

10 Display the VLANs configured on STP-enabled ports. For example:

```
admin> stp -v
VlanId  Bridge-Group
-----  -
      10      10
```

This command displays VLAN 10 and the bridge group to which it belongs.

Note that the LAN “router” interface defined in the BRIDGE-GROUP profile will be supported on both STP-enabled ports.

Routing implications for STP-enabled ports

STP is a layer 2 control protocol for bridged traffic. However, the Gigabit Ethernet ports of the IP2100 control module handle both bridged and routed traffic, and STP state changes on the ports affect routed traffic as well as bridged traffic. If not properly handled, STP state changes on the ports would trigger a flood of route update messages and the network might take time to converge, which could block routed traffic for a long time.

The Stinger RSTP implementation ensures that state changes of the STP-enabled ports will not trigger route update messages and hence the network or routing protocols do not attempt to converge. In fact, there will be no change to the routing table and routed traffic will not be blocked for a long time.

To prevent route update messages when a Gigabit Ethernet port transitions from a Blocking to a Forwarding state, both ports must keep track of a next hop router’s reachability. To enable this, the system monitors downstream traffic and constructs a table that maps the combination of VLAN ID and MAC address to specific ports. In

addition, the system sends ARP requests whenever the STP state of one of the ports changes.

With these two actions, nothing changes from the routing perspective on STP-enabled ports, and no route update messages are generated when one port switches to a Blocking state. The following command displays the contents of the VLAN-MAC to port mapping table:

```
admin> macStpportMap
  MAC Address      VlanID IfNum
-----
00:00:e2:19:bd:16  2111   33
```

The first column of the output displays the MAC address of the gateway. The second column displays the VLAN ID associated with the MAC address in the first column. The third column displays the interface number of the Gigabit Ethernet port on which the MAC frame was received.

Administrative tools for Gigabit Ethernet

The system supports the `gmac` command for administrative information about Gigabit Ethernet ports. If you are managing the system remotely, some of this information is also available through the `ip2kstats` MIB. For details, see “Gigabit Ethernet diagnostics” on page A-2 and “SNMP MIB for GMAC and VLAN statistics” on page A-29. For other commands that can be used to monitor activity on any Ethernet port, such as `etherdisplay`, see the *Stinger Reference*.

VLAN Configuration

3

IP filters for Ethernet-encapsulated bridged IP datagrams	3-2
Configuring 1:1 VLAN bridging	3-2
Configuring N:1 VLAN bridging	3-6
Configuring stacked VLANs	3-27
Configuring routed VLANs	3-35
VLAN bridging of IPoA traffic to an upstream BRAS	3-39
PPPoA bridging in 1:1 and N:1 VLAN	3-49
Administrative tools for VLAN	3-56

This chapter describes configuration of virtual LAN (VLAN) operations and traffic characteristics. For detailed background information about VLAN, see *IEEE standard 802.1Q (1998) for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks*. In this chapter, the following VLAN terminology is used:

Table 3-1. Definition of VLAN terms

VLAN term	Definition
VLAN bridging (1:1)	A bridging configuration between a single PVC and a VLAN, with a 1:1 mapping. The setup uses <code>vlan-circuit</code> bridging.
VLAN bridging (N:1)	A bridging configuration between multiple PVCs and a VLAN, using <code>transparent</code> bridging and a bridge group.
Stacked VLAN	A bridging configuration that encapsulates a VLAN within another VLAN to greatly increase the VLAN space. A stacked VLAN uses double tagging, where one 802.1Q tag represents a service provider (NSP) and a second 802.1Q tag represents an ID that is unique within the NSP VLAN.
Routed VLAN	A routing configuration in which the upstream IP interface is VLAN enabled. A routed VLAN interface is always mapped to a virtual IP interface on a control module GigE port. Packets received on a routed VLAN interface are routed based on the IP address, and packets are sent through the routed VLAN interface based on an IP routing decision.

IP filters for Ethernet-encapsulated bridged IP datagrams

Systems with an IP2100 control module support IP filtering applied to Ethernet-encapsulated bridged IP datagrams on the DSL side of the following types of Layer 2 connections:

- 1:1 bridged VLAN circuits
- N:1 bridged VLAN circuits (with transparent bridging)
- Stacked VLAN circuits

In terms of defining and verifying the IP filters, all configurations and commands remain the same as for IP filters applied to routed connections. See Chapter 11, “Filter Configuration.”

The IP filters apply only to CONNECTION profiles. Layer 3 filters are not supported on bridged Ethernet interfaces. Although there is no runtime validation to prevent you from applying a Layer 3 filter in a VLAN-ETHERNET profile, it is a misconfiguration and would not have the desired effect.


The following restrictions apply:

- IP filters apply *only* to Ethernet-encapsulated bridged IP datagrams. The filters do not apply to IP packets with other types of encapsulations, such as PPPoE.
- This feature is supported only on systems with an IP2100 control module.

Configuring 1:1 VLAN bridging

This section describes how to configure `vlan-circuit` bridging, which maps one user PVC to one unique VLAN ID, as shown in Figure 3-1. Source MAC address learning does not apply (and is not needed) in this configuration.

Figure 3-1. Bridging VLAN: One PVC to one VLAN (1:1)



With the VLAN circuit (1:1) bridging configuration, the user CPE encapsulates data using RFC 2684 encapsulation for bridged protocols. The system bridges the frames received on the user PVC only to the paired VLAN, and vice versa.

To configure `vlan-circuit` bridging, complete the following steps:

- 1 Create a VLAN-ETHERNET profile on Gigabit Ethernet and assign a bridge group number.
- 2 Create a CONNECTION profile and specify the same bridge group.

Overview of VLAN-ETHERNET and CONNECTION settings

The index of a VLAN-ETHERNET profile specifies the physical address of the Gigabit Ethernet port and a unique VLAN ID. Following are the profile contents, shown for VLAN ID 50:

```
[in VLAN-ETHERNET/{ { shelf-1 first-control-module 2 } 50 }]
interface-address* = { { shelf-1 first-control-module 2 } 50 }
vlan-id = 0
enabled = no
filter-name = ""
pppoe-options = { no no }
bridging-options = { 0 no no vlan-circuit 0 }

[in VLAN-ETHERNET/{ { shelf-1 first-control-module 2 } 50}:bridging-options]
bridging-group = 0
bridge = no
bridge-type = vlan-circuit
mac-address-learning-limit = 16
```

The following connection settings, shown with default settings, enable subscribers to bridge to the configured VLAN:

```
[in CONNECTION/"":bridging-options]
bridging-group = 0
bridge = no
bridge-type = vlan-circuit
mac-address-learning-limit = 16
```

Parameter	Setting
interface-address	Address of the Gigabit Ethernet port followed by the VLAN ID, using the following format: <pre>{ { shelf-n slot-n port-n } vlan-id }</pre> The slot-n is first-control-module or second-control-module, depending on the slot in which the active control module is installed, and port-n is 2 for the Gigabit Ethernet port. The vlan-id value is the IEEE 802.1Q VLAN tag value added to the IP packets transmitted on the Ethernet interface. The valid range is from 0 to 4095, but for full compatibility with IEEE 802.1Q, Alcatel-Lucent recommends that you do not use the vlan-id values of 0, 1 or 4095. However, the system does not prevent you from assigning these values.
vlan-id	VLAN ID. This setting is read-only. You must set it in the index of the VLAN-ETHERNET profile.
enabled	Enable/disable the VLAN-ETHERNET profile.


Parameter	Setting								
filter-name	Name of a filter to apply to traffic received on the VLAN interface. For all types of bridged VLANs, the configured filter must be an Ethernet input filter. For routed VLANs, the filter must be an IP filter. For details about creating and applying filters, see Chapter 11, “Filter Configuration.”								
pppoe-options	<i>Not currently supported.</i>								
bridging-group	Number from 0 to 65535, used to group bridged interfaces. For 1:1 VLAN bridging, this setting must match in the VLAN-ETHERNET and CONNECTION profiles. For N:1 VLAN bridging, this setting must match in the BRIDGE-GROUP, VLAN-ETHERNET and CONNECTION profiles.								
bridge	If enabled, associates the port with the specified bridging-group.								
bridge-type	Type of bridging to use on the VLAN or WAN interface: <table> <tr> <td>vlan-circuit</td><td>Required setting for 1:1 VLAN bridging (the default).</td></tr> <tr> <td>transparent-bridging</td><td>Required setting for N:1 VLAN bridging.</td></tr> <tr> <td>stacked-vlan</td><td>Required setting for stacked VLAN bridging.</td></tr> <tr> <td>no-bridging</td><td>Required setting for routed VLANs.</td></tr> </table>	vlan-circuit	Required setting for 1:1 VLAN bridging (the default).	transparent-bridging	Required setting for N:1 VLAN bridging.	stacked-vlan	Required setting for stacked VLAN bridging.	no-bridging	Required setting for routed VLANs.
vlan-circuit	Required setting for 1:1 VLAN bridging (the default).								
transparent-bridging	Required setting for N:1 VLAN bridging.								
stacked-vlan	Required setting for stacked VLAN bridging.								
no-bridging	Required setting for routed VLANs.								
mac-address-learning-limit	Maximum number of MAC addresses the system will learn across the interface. With a zero (0) setting, no limit is set in software. This is the default setting on a VLAN interface. With a nonzero value, the system adds only the specified number of addresses to the bridge table. The maximum number of MAC addresses the system will learn on a DSL or VLAN interface is 1024. In a CONNECTION profile, a new mac-address-learning-limit setting takes effect when you write the profile. The connection is bounced to instantiate the new limit. However, for a new value to take effect on the VLAN Ethernet side, you must manually disable the profile and then reenab le it. The bridge-type parameter must be set to transparent-bridging for this setting to take effect. For more details, see “How address limiting works” on page 3-11.								

Parameter	Setting
packet-type	Type of packet to be bridged. With the default none setting, all packet types are bridged. Set <code>packet-type</code> to <code>ipoa</code> to bridge IPoA packets on this connection and responds to ARP requests received on the VLAN Ethernet interface. The <code>bridge-type</code> parameter must be set to <code>vlan-circuit</code> (the default) for this setting to take effect.

Sample 1:1 VLAN bridging configuration

Figure 3-2 shows a Stinger system bridging a PVC to a VLAN:

Figure 3-2. Sample 1:1 VLAN circuit



To configure the VLAN circuit, first verify that bridging is enabled on the physical interface. For example:

```
admin> get ethernet { 1 8 2 } bridging-enabled
[in ETHERNET/{ shelf-1 first-control-module 2}:bridging-enabled]
bridging-enabled = yes
```

If bridging is not enabled, enable it as described in “Enabling layer 2 bridging for VLAN operations” on page 2-3. Then, follow these steps:

- 1 Create a VLAN-ETHERNET profile.

```
admin> new vlan-ethernet { { 1 8 2 } 50 }
admin> set enabled = yes
admin> set bridging-options bridging-group = 34590
admin> write -f
```
- 2 Create a CONNECTION profile for the PVC side of the VLAN circuit.

```
admin> new connection dslcpe
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options ip-routing-enabled = no
admin> set bridging-options bridging-group = 34590
admin> set bridging-options bridge = yes
admin> set atm-options vpi = 8
admin> set atm-options vci = 100
```


```
admin> set atm-options nailed-group = 51
admin> write -f
```

For background information about configuring PVCs, see the *Stinger ATM Configuration Guide*.

Configuring N:1 VLAN bridging

This section describes how to bridge multiple user PVCs onto a VLAN, as shown in Figure 3-3.

Figure 3-3. Bridging multiple PVCs to a VLAN



The user CPE encapsulates data using RFC 2684 encapsulation for bridged protocols, and the system bridges the frames to the associated VLAN (and vice versa), performing transparent bridging to build a table of known MAC addresses and the port associated with each address. If the system receives packets for an unknown MAC address, or if it receives broadcast packets, the traffic is forwarded on all ports that are part of the bridge group.

A BRIDGE-GROUP profile defines the traffic characteristics of the VLAN and assigns a BRIDGE-GROUP number that is shared by the VLAN configuration and all user PVCs that are part of it.

To configure VLAN bridging, complete the following steps:

- 1 Create a BRIDGE-GROUP profile that defines traffic characteristics for the VLAN and assigns its bridging-group number.
- 2 Create a VLAN-ETHERNET profile on the Gigabit Ethernet port and bind it to the BRIDGE-GROUP profile.
- 3 Create CONNECTION profiles for user PVCs, and bind them to the same bridge group.

Creating and configuring bridge groups

To define a limited broadcast domain for a bridged VLAN with multiple subscriber interfaces, the system requires bridge groups. Each VLAN must have a unique bridge group number, which is shared by all interfaces within the same VLAN.



Note The maximum number of bridge groups in the system is 256. However, the practical limit will be less than that if multicast groups are also configured. For details, see “IP2100 and IP2000 maximum limits on multicast and bridge groups” on page 9-3.

The system creates a bridging table that contains only interfaces in the same bridge group, and when it receives packets from one interface in the group, it consults only

that bridging table for destination ports. It will not forward the traffic to interfaces that are not in the same bridge group.

To optimize its forwarding operations over time, the system uses an IEEE 802.1 transparent bridging algorithm to build a table of known MAC addresses and the port associated with each address. If it receives packets for an unknown MAC address, or if it receives broadcast packets, the traffic is forwarded on all ports that are part of the bridge group except the port on which the packets were received.

Overview of BRIDGE-GROUP settings

The BRIDGE-GROUP profile defines traffic characteristics for VLAN bridging. The index of a BRIDGE-GROUP profile specifies a **bridging-group** number. Following are the BRIDGE-GROUP parameters, shown with default settings:

```
[in BRIDGE-GROUP/0]
enable = no
bridging-group = 0
mac-entry-age-time = 300
igmp-snooping = no
port-block-enabled = yes
lan-router-interface-address = { { any-shelf any-slot 0 } 0 }
wan-router-interface-profile = ""
dhcp-snooping = { no { no 0.0.0.0 "" 1 no } { no 0.0.0.0 1 "" no } }
pppoe-snooping = { no { no 0.0.0.0 "" 1 no } { no 0.0.0.0 "" 1 no } }
```

Parameter	Setting
enable	Enables or disables the BRIDGE-GROUP profile. Set this parameter to yes .
bridging-group	Number from 0 to 65535, used to group bridged interfaces. The same number specified here must also be specified as the bridging-group number in the VLAN-ETHERNET and CONNECTION profiles.
mac-entry-age-time	Timeout interval (in seconds) at which the system ages out inactive MAC addresses from the bridge group's bridging table. Valid values are from 0 (which disables the address aging function) to 65535. The default is 300 seconds. When this timer expires, the system traverses all source MAC addresses in the group's bridge table and deletes those addresses from which no traffic has been received since the last traversal. A maximum of 32K addresses can be aged out systemwide. If you modify this value for a bridge group that is in use, the new value is used for the next timeout.

Parameter	Setting
igmp-snooping	Enables or disables IGMP snooping. When IGMP is disabled (the default), multicast data streams are forwarded to all ports in the VLAN, even those who have not registered for the multicast. When IGMP snooping is enabled and a Join is received from a subscriber interface, the system snoops the packet and makes an entry in its bridge table, along with the IP multicast address. This Join is forwarded only on the “router” interface and not on the other interfaces in the bridge group. To use IGMP snooping, set this parameter to yes, and configure the <code>lan-router-interface-address</code> or <code>wan-router-interface-address</code> parameter.
port-block-enabled	Enables or disables port blocking to prevent traffic flows between subscriber interfaces in the same VLAN. Port blocking is enabled by default, but to use it you must configure the <code>lan-router-interface-address</code> or <code>wan-router-interface-address</code> parameter.
lan-router-interface-address	Index of the VLAN-ETHERNET profile to be used as the “router” interface in terms of traffic handling for this bridge group. The specified interface is the default path for uplinks from users, and is used in port blocking to prevent other, unintended uses of the subscriber PVC interfaces. If IGMP snooping is enabled, the index must specify the Ethernet or VLAN interface on which downstream multicast streams are received. This setting and the <code>wan-router-interface-address</code> setting are mutually exclusive.
wan-router-interface-profile	Name of the CONNECTION profile to be used as the “router” port in terms of traffic handling for this bridge group. The specified connection must be on a trunk interface. It is the default path for uplinks from users, and is used in port blocking to prevent other, unintended uses of the subscriber PVC interfaces. If IGMP snooping is enabled, the CONNECTION profile name must specify the trunk interface on which downstream multicast streams are received. This setting and the <code>lan-router-interface-address</code> setting are mutually exclusive.

Parameter	Setting
dhcp-snooping	Subprofiles for enabling DHCP snooping or PPPoE snooping (or both). When DHCP and PPPoE snooping are enabled, the system snoops clients' DHCP requests and PPPoE Discovery packets and add identifying information to the packets before bridging the packets upstream. The unique identifiers can be used in authentication and accounting, and for troubleshooting purposes. For details, see "Configuring DHCP and PPPoE snooping for DSL line identification" on page 3-14.
pppoe-snooping	

Sample BRIDGE-GROUP configuration with MAC address aging

The following commands modify bridge-group 451 to extend the duration of the timer for discovering and dropping inactive source MAC addresses from its bridge tables from the default five minutes (300 seconds) to 10 minutes.

```
admin> read bridge-group 451
admin> set mac-entry-age-time = 600
admin> write -f
```

Sample BRIDGE-GROUP configuration with port blocking

Port blocking prevents users who are bridged to the same LAN or VLAN from exchanging traffic flows user-to-user or from building virtual networks. To use it, you must identify one "router" interface for the bridge group, which can be a VLAN interface or a CONNECTION profile on a trunk interface. All other ports in the same bridge group are then considered subscriber interfaces for traffic handling purposes. Traffic received on the "router" and subscriber interfaces in a bridge group is restricted as shown in Table 3-2.

Table 3-2. Traffic restrictions when port blocking is enabled

Received on "router" port	Received on subscriber ports
No restrictions are placed on traffic received on the "router" interface.	Packets received on a DSL port and destined for another DSL port in the same bridge group are discarded. This prevents users from building virtual networks using the VLAN, or for making user-to-user connections through the VLAN.
	Broadcast packets are forwarded only to the "router" interface.
	Packets destined for a MAC address the system does not recognize are forwarded only to the "router" interface.

The next commands modify bridge-group 275 (VLAN 500) to use DSL port blocking, by specifying a “router” interface:

```
admin> read bridge-group 275
admin> set lan-router-interface physical shelf = shelf-1
admin> set lan-router-interface physical slot = first-control-module
admin> set lan-router-interface physical item = 2
admin> set lan-router-interface logical-item = 500
admin> write -f
```

Sample BRIDGE-GROUP configuration with IGMP snooping

Layer 2 multicasting maps an IP multicast traffic to a MAC multicast address, which is treated at layer 2 as broadcast traffic. Because broadcast traffic would typically be forwarded to all interfaces in the VLAN, IGMP snooping is implemented on a bridge group basis to enable efficient support for layer 2 multicasting to VLANs.

IGMP snooping “peeks” into the layer 3 content of multicast packets, and allows the IP control module to forward multicast traffic for a particular group only to those user PVCs that have registered in the group.

When IGMP snooping is enabled within a bridge group, all IGMP packets received on an interface in that bridge group are snooped, and multicast forwarding is done on the basis of the multicast group address and not on the basis of a multicast MAC address.

Multicast data traffic is forwarded only to the users subscribed to the particular multicast group. Table 3-3 shows how IGMP control packets are handled.

Table 3-3. IGMP control packet handling with IGMP snooping enabled

Control packet type	System action
IGMP-QUERY	Generic queries received from the “router” interface are forwarded to all DSL users.
	Group-specific queries received from the “router” interface are forwarded to users who have joined that group.
IGMP-REPORT(v1/v2)	Reports received from users are forwarded only to the “router” interface.



Note The Stinger system does not generate any IGMP queries or reports. It is up to the upstream router to generate queries and handle the reports, and it is up to the end users to send reports.

The following commands modify bridge-group 22 (VLAN 478) to enable IGMP snooping. A designated “router” interface is required for this feature to work:

```
admin> read bridge-group 22
admin> set igmp-snooping = yes
admin> set lan-router-interface physical shelf = shelf-1
admin> set lan-router-interface physical slot = first-control-module
```

```
admin> set lan-router-interface physical item = 2
admin> set lan-router-interface logical-item = 478
admin> write -f
```

When IGMP snooping is enabled and a Join is received from a subscriber interface, the system snoops the packet and makes an entry in its bridge table, along with the IP multicast address. This Join is forwarded only on the “router” interface and not on the other interfaces in the bridge group.

Similarly, when a multicast data packet arrives from the LAN router interface, it is snooped, checked for the IP multicast address and forwarded only on subscriber interfaces from which a Join was received.

If both IGMP snooping and port blocking are enabled and the system receives traffic for 224.0.0.* from the “router” interface, it is forwarded on all the ports in the bridge group. However, if traffic for 224.0.0.* is received from a subscriber interface, it is forwarded only to the “router” interface.

VLAN and connection settings

For details about the VLAN and bridging parameters used in the `vlan-ethernet` and `CONNECTION` profiles for an N:1 VLAN bridging configuration, see “Overview of VLAN-ETHERNET and CONNECTION settings” on page 3-3. For background information about configuring subscriber PVCs, see the *Stinger ATM Configuration Guide*.

How address limiting works

The transparent bridging algorithm enables the system to learn the MAC addresses of devices sending traffic across an interface and enter the addresses in an internal bridge table. This algorithm is subject to the limits shown in Table 3-4, which are based on the amount of memory reserved for this purpose.

Table 3-4. Maximum number of learned MAC addresses per module

Module	Types of interfaces	Total number of learned addresses
Control module	Trunk or GigE	32,768
LIM	DSL	1,170

Within the limits shown in Table 3-4, you can restrict the number of addresses the system will learn across a connection by setting the `mac-address-learning-limit` parameter in the `connection bridging-options` subprofile. If the system receives traffic from an unknown source address after learning the maximum number of MAC addresses, it discards the traffic and does not add the source address to the bridge table.

For example, with the default 16 MAC address learning limit per connection, if 72 bridging connections are brought up on a LIM, the number of learned addresses ($16 * 72$) is 1,152, which is within the 1,170 limit on that LIM. If you require more than 72 bridging calls per LIM, you must correspondingly reduce the MAC address learning limit per connection to remain within the LIM-wide limit of 1170.


Limiting the number of source MAC addresses learned on a transparent bridging interface restricts the number of users that can access the network through a single CPE, and prevents a type of denial-of-service attack in which a user overloads the bridge table by sending heavy traffic from many different source MAC addresses. When the address limit has been reached on a connection, the system logs a message to that effect, to enable the administrator to check for denial-of-service attacks or, if the connection attempts are valid, to reconfigure the limit on a connection.

If MAC address aging is also enabled, when one of the devices on an interface becomes inactive for a specified interval, the system ages the address out of its bridge table. As soon as the bridge table contains less than the maximum number of addresses for the interface, the system can again accept traffic from an unknown source MAC address on that interface and add the address to its bridge table. For details about address aging, see “Sample BRIDGE-GROUP configuration with MAC address aging” on page 3-9.

Sample N:1 VLAN bridging configuration with address limiting

Figure 3-4 shows a Stinger system bridging multiple user PVCs onto a VLAN. MAC address limiting will be enforced for CPE-1.

Figure 3-4. Sample N:1 VLAN bridging



With the sample configuration shown below, the system learns up to three MAC addresses on the CPE-1 connection. For the first three PCs that send traffic, the system forwards the traffic and learns the source MAC addresses. If the fourth PC attempts to send traffic, the user's connection is rejected. Later, if one of the bridged PCs does not send traffic for a duration equal to the `mac-entry-age-time` setting, the system ages out that MAC address from its bridge table. If the fourth PC sends traffic at that time, the system learns its address and forwards the traffic.

To configure this N:1 VLAN bridging setup, first verify that bridging is enabled on the physical interface. For example:

```
admin> get ethernet { 1 8 2 } bridging-enabled
[in ETHERNET/{ shelf-1 first-control-module 2}:bridging-enabled]
bridging-enabled = yes
```

If bridging is not enabled, enable it as described in “Enabling layer 2 bridging for VLAN operations” on page 2-3. Then, follow these steps:

- 1 Create a bridge group. The following group sets the address age-out interval to three minutes and uses DSL port blocking. (For more detail, see “Creating and configuring bridge groups” on page 3-6.)

```
admin> new bridge-group
admin> set enable = yes
admin> set bridging-group = 451

admin> set lan-router-interface physical shelf = shelf-1
admin> set lan-router-interface physical slot = first-control-module
admin> set lan-router-interface physical item = 2
admin> set lan-router-interface logical-item = 50
admin> set mac-entry-age-time = 180
admin> write -f
```

- 2 Create a VLAN-ETHERNET profile.

```
admin> new vlan-ethernet
admin> set interface-address physical shelf = shelf-1
admin> set interface-address physical slot = first-control-module
admin> set interface-address physical item-number = 2
admin> set interface-address logical-item = 50
admin> set enabled = yes
admin> set bridging-options bridging-group = 451
admin> set bridging-options bridge = yes
admin> set bridging-options bridge-type = transparent-bridging
admin> write -f
```

For an example of applying a filter to a bridged VLAN interface, see “Applying a filter to a VLAN Ethernet interface” on page 11-22.

- 3 Create CONNECTION profiles for bridged PVCs to the CPE devices. The profiles must specify the right bridge group number. The profile for CPE-1 also limits the number of MAC addresses to three.

```
admin> new connection cpe-1
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options ip-routing = no
admin> set bridging-options bridge = yes
admin> set bridging-options bridging-group = 451
admin> set bridging-options bridge-type = transparent-bridging
admin> set bridging-options mac-address-learning-limit = 3
admin> set atm-options vpi = 0
```

```
admin> set atm-options vci = 25
admin> set atm-options nailed-group = 51
admin> write -f
admin> new connection cpe-2
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options ip-routing = no
admin> set bridging-options bridge = yes
admin> set bridging-options bridging-group = 451
admin> set bridging-options bridge-type = transparent-bridging
admin> set atm-options vpi = 0
admin> set atm-options vci = 38
admin> set atm-options nailed-group = 57
admin> write -f
```

Configuring DHCP and PPPoE snooping for DSL line identification

You can configure DHCP snooping or PPPoE snooping, or both, to bridge unique DSL line identification to upstream network elements. The unique identifier can be used in authentication and accounting, and for troubleshooting purposes.

When DHCP and PPPoE snooping are enabled, the system snoops clients' DHCP requests and PPPoE Discovery packets on transparent bridging connections and adds identifying information to the packets before forwarding the packets upstream. Similarly, the system snoops downstream replies from a server and removes the identifiers before bridging the packets to the DSL client.

This functionality is configurable by bridge group. The interface on which the packets are forwarded upstream is the "router" interface defined in the BRIDGE-GROUP profile, which means that the modified packets can be forwarded upstream on a trunk interface (wan-router-interface) or Ethernet (lan-router-interface), but not both.


Structure of identifiers added to DHCP and PPPoE Discovery packets

This section describes the identifying information added to DHCP Request and PPPoE Discovery packet headers before they are forwarded upstream when DHCP or PPPoE snooping is enabled.

Including a Stinger interface IP address

Upstream network elements can use the `version` and `if-ip` fields (shown in Figure 3-5) to identify the Stinger system. These values are explicitly configured in the BRIDGE-GROUP profile.

Figure 3-5. Format when an interface IP address (**if-ip**) is used



The **hostname** field, which is automatically encoded from the hostname of the CONNECTION profile, identifies the DSL end connection.

Including an identifying text string (PPPoE snooping only)

Because snooping functionality occurs at layer 2 (so the IP address of the Stinger system is not directly relevant), the **if-ip** field can optionally be replaced by a text string used to identify the Stinger system. In that case, upstream network elements use the **version** and **vendor-option-string** fields (shown in Figure 3-6) to identify the Stinger system. These values are explicitly configured in the BRIDGE-GROUP profile.

Figure 3-6. Format when a text string (**vendor-option-string**) is used



The **hostname** field, which is automatically encoded from the hostname of the CONNECTION profile, identifies the DSL end connection.




Note Currently, this option is supported only for PPPoE snooping. The **vendor-option-string** and **if-ip** fields are mutually exclusive.

Including only the hostname identifier

Typically the system adds the entire suboption field, including the **hostname** and the **version**, **if-ip**, and other fields, using default values if the fields are unconfigured.

You can specify that only the **hostname** should be sent in the suboption field, as shown in Figure 3-7. Only the actual **hostname** and its length is sent, not the maximum allowed 32-byte field. Other values within the subprofile (such as **version** or **if-ip**) are not sent, even if they are configured.

Figure 3-7. Format when sending only the hostname (**send-only-hostname**)



Configuring DHCP snooping

With DHCP snooping, the Stinger system snoops client-to-server DHCP request packets and add identifiers before bridging the packets upstream. Source MAC address learning is performed for packets that are forwarded via transparent bridging.

The system also snoops server-to-client DHCP reply packets to remove identifiers before bridging the packets downstream to the destination MAC address of the frame. If it receives DHCP reply packets from the server that do not contain the option 82 identifiers and DHCP snooping is enabled, the system drops the packets.

DHCP option 82 has two suboptions:

- Suboption 1, for conveying a unique identifier for the user circuit
- Suboption 2, for conveying a unique identifier for the remote host

You can configure the system to convey both suboptions or only one of them. For DHCP snooping, the two suboptions behave similarly in that both suboptions encode the hostname of the CONNECTION profile to identify the DSL end connection. For details about the format of the identifiers added by DHCP snooping, see “Structure of identifiers added to DHCP and PPPoE Discovery packets” on page 3-14.

Overview of DHCP-snooping settings in BRIDGE-GROUP profiles

Following are the subprofiles and parameters, shown with default settings, for configuring DHCP snooping for a bridge group. Configure suboption 1 in the `circuit-id` subprofile or suboption 2 in the `remote-id` subprofile, or both. For DHCP snooping, both suboptions encode the hostname of the user VC CONNECTION profile.

```
[in BRIDGE-GROUP/0:dhcp-snooping]
enable = no
circuit-id = { no no 0.0.0.0 "" 1 no }
remote-id = { no no 0.0.0.0 "" 1 no }

[in BRIDGE-GROUP/0:dhcp-snooping:circuit-id ]
enable = no
send-hostname-only = no
if-ip = 0.0.0.0
vendor-option-string = ""
version = 1
dhcp-allow-any-src-port = no


[in BRIDGE-GROUP/0:dhcp-snooping:remote-id]
enable = no
send-hostname-only = no
if-ip = 0.0.0.0
vendor-option-string = ""
version = 1
dhcp-allow-any-src-port = no
```


Parameter	Description
enable	Enable or disable DHCP snooping for this bridge group.
circuit-id remote-id: enable	Enables or disables suboption 1 (in the circuit-id subprofile) or suboption 2 (in the remote-id subprofile). If enabled, the system encodes the hostname of the PVC on which the DHCP Request packet was received.
circuit-id remote-id: send-hostname-only	Setting this parameter to yes causes the system to send only the hostname and its length in the DHCP Option 82 suboption field. See “Including only the hostname identifier” on page 3-15.
circuit-id remote-id: if-ip	IP address used to identify the Stinger system to the DHCP server. If both circuit and remote ID options are enabled, only one address is used. If both this field and vendor-option-string are empty, the Stinger uses the system address if that value has been defined. To interoperate with DHCP servers that zero-delimit suboption fields, the specified IP address cannot contain a zero octet. This setting and the vendor-option-string setting are mutually exclusive in the same subprofile.
circuit-id remote-id: vendor-option-string	Currently, this option is supported only for PPPoE snooping. For details, see “Including an identifying text string (PPPoE snooping only)” on page 3-15.
circuit-id remote-id: version	Configurable version ID, set to 1 by default. To interoperate with DHCP servers that zero-delimit suboption fields, change this setting to a value of 257 or higher.
circuit-id remote-id: dhcp-allow-any-src-port	Currently, this option is supported only for IPoA and BIR connections. For details, see “Interoperation with DHCP servers that zero-delimit suboption fields” on page 4-76.

Sample DHCP snooping configuration

Figure 3-8 shows a Stinger system bridging user PVCs onto an Ethernet VLAN using transparent bridging. The CPE is operating in bridging mode, and the DHCP clients obtain their configuration from an upstream DHCP server.

Figure 3-8. DHCP snooping example



To configure this VLAN for DHCP snooping, first verify that bridging is enabled on the physical interface. For example:

```
admin> get ethernet { 1 8 2 } bridging-enabled
[in ETHERNET/{ shelf-1 first-control-module 2}:bridging-enabled]
bridging-enabled = yes
```

Then, follow these steps on the Stinger system:

- 1 Create a bridge group and enable DHCP snooping. You must also configure the `lan-router` interface and configure the `remote-id` or `circuit-id` subprofile, or both. In this example, the `circuit-id` subprofile is configured, and the Stinger system's IP address is used to identify it as the intermediate agent.

```
admin> new bridge-group
admin> set enable = yes
admin> set bridging-group = 333
admin> set lan-router-interface physical shelf = shelf-1
admin> set lan-router-interface physical slot = first-control-module
admin> set lan-router-interface physical item = 2
admin> set lan-router-interface logical-item = 18
admin> set dhcp-snooping enable = yes
admin> set dhcp-snooping circuit-id enable = yes
admin> set dhcp-snooping circuit-id if-ip = 8.8.8.8
admin> write -f
```

- 2 Create a VLAN-ETHERNET profile.

```
admin> new vlan-ethernet
admin> set interface-address physical shelf = shelf-1
admin> set interface-address physical slot = first-control-module
admin> set interface-address physical item = 2
admin> set interface-address logical-item = 18
admin> set enabled = yes
admin> set bridging-options bridging-group = 333
admin> set bridging-options bridge = yes
admin> set bridging-options bridge-type = transparent-bridging
admin> write -f
```


- 3 Create CONNECTION profiles for bridged PVCs to the CPE device. The profiles must specify the right bridge group number.

```
admin> new connection
admin> set station = dhcp-client-1
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options ip-routing-enabled = no
admin> set bridging-options bridging-group = 333
admin> set bridging-options bridge = yes
```

```
admin> set bridging-options bridge-type = transparent-bridging
admin> set atm-options vpi = 0
admin> set atm-options vci = 57
admin> set atm-options nailed-group = 304
admin> write -f
```

With this configuration, when the client generates a DHCP request, the system adds the following fields to the request packet:

Figure 3-9. Contents of fields with sample DHCP snooping configuration



When the DHCP server replies with options, the Stinger system removes those options before forwarding the reply to the client based on the destination MAC address in the frame. If it receives DHCP replies from the server without option 82, the system drops the packets.

Sample bridge-group configuration for sending only a hostname

The next commands configure a bridge group for a VLAN interface on which only the hostname will be added to upstream packets:

```
admin> new bridge-group
admin> set enable = yes
admin> set bridging-group = 500
admin> set port-block-enabled = no
admin> set lan-router-interface physical-address shelf = shelf-1
admin> set lan-router-interface physical-address slot = first-control-module
admin> set lan-router-interface physical-address item-number = 2
admin> set lan-router-interface logical-item = 500
admin> set dhcp-snooping enable = yes
admin> set dhcp-snooping circuit-id enable = yes
admin> set dhcp-snooping circuit-id send-hostname-only = yes
admin> write -f
```

With this configuration, when the client in this bridge group generates a DHCP request, the system adds only the number of octets in the hostname and the hostname of the PVC on which the DHCP Request was received.

Configuring PPPoE snooping for line identification of PPPoE clients

When PPPoE snooping has been configured for the bridge group, the Stinger system snoops incoming PPPoE Discover packets (Ether Type 0x8863) on connections and adds vendor-specific information before bridging packets to an upstream interface such as Ethernet.

For downstream PPPoE packets (server-to-client), the system snoops the packets to remove the vendor-specific options, and then bridges the packet based on the destination MAC address of the frame. Source MAC address learning is performed for all packets that are forwarded via transparent bridging.

PPPoE Discovery packet handling with PPPoE snooping

When PPPoE snooping is enabled, the system snoops the following types of PPPoE Discovery packets:

- PPPoE Active Discovery Initiation (PADI)
- PPPoE Active Discovery Offer (PADO)
- PPPoE Active Discovery Request (PADR)
- PPPoE Active Discovery Session-confirmation (PADS)
- PPPoE Active Discovery Terminate (PADT)

Table 3-5 shows how these packets are handled in the upstream and downstream directions.

Table 3-5. Packet handling with PPPoE snooping

Traffic direction	PPPoE Discovery packet handling
Upstream	For PPPoE Discovery packets received on a client interface within the bridge group (any interface other than the “router” interface), the following actions are performed: ■ PADO and PADS packets are discarded. ■ PADT packets are forwarded without modification via transparent bridging. ■ PADI and PADR packets are modified by adding vendor-specific options and are then forwarded via transparent bridging. Because PADI is a broadcast packet, it is forwarded to all ports in the bridge group if DSL port blocking is not enabled. If a packet’s length exceeds 1500 after adding the vendor-specific options, the Stinger system discards the packet. If PADI or PADR packets coming from the clients already contain the vendor-specific options, the Stinger system discards the packet.

Table 3-5. Packet handling with PPPoE snooping (Continued)

Traffic direction	PPPoE Discovery packet handling
Downstream	For PPPoE Discovery packets received on the “router” interface of the bridge group (as defined by the <code>lan-router-interface</code> or <code>wan-router-interface</code> setting), the following actions are performed: ■ PADI and PADR packets are discarded. ■ PADT packets are forwarded without modification via transparent bridging. ■ PADO and PADS packets are modified by removing vendor-specific options if any are present, and are then forwarded via transparent bridging.

Vendor-specific options in PPPoE packets

The vendor-specific tag added to or removed from PPPoE Discovery packets contains the fields shown in Figure 3-10.

Figure 3-10. PPPoE vendor-specific tag format

0x0105 (Vendor-Specific)	TAG_LENGTH
0x00000DE9 (ADSL Forum IANA entry)	
0x01	length
Agent Circuit ID value (cont)	Agent Circuit ID value...
0x02	length
Agent Remote ID value (cont)	Agent Remote ID value...

The 0x0105 field is the vendor-specific tag described in RFC 2516.

The 0x01 and 0x02 fields contain DHCP option 82 suboption values (circuit-id and remote-id)

The 0x01 and 0x02 fields contain values from the `pppoe-snooping` configuration in the `BRIDGE-GROUP` profile, as described in “Structure of identifiers added to DHCP and PPPoE Discovery packets” on page 3-14.

Overview of PPPoE snooping settings in BRIDGE-GROUP profiles

Following are the subprofiles and parameters, shown with default settings, for configuring PPPoE snooping:

```
[in BRIDGE-GROUP/0:pppoe-snooping]
enable = no
circuit-id = { no no 0.0.0.0 "" 1 no }
remote-id = { no no 0.0.0.0 "" 1 no }

[in BRIDGE-GROUP/0:pppoe-snooping:circuit-id]
enable = no
send-hostname-only = no
if-ip = 0.0.0.0
```

```

vendor-option-string = ""
version = 1
dhcp-allow-any-src-port = no
[in BRIDGE-GROUP/0:pppoe-snooping:remote-id]
enable = no
send-hostname-only = no
if-ip = 0.0.0.0
vendor-option-string = ""
version = 1
dhcp-allow-any-src-port = no


```

Parameter	Description
enable	Enable or disable PPPOE snooping for this bridge group.
circuit-id remote-id: enable	Enables or disables suboption 1 (in the circuit-id subprofile) or suboption 2 (in the remote-id subprofile). If enabled, the system encodes the station value (the hostname) of the CONNECTION profile that defines the PVC on which the PPPoE Discovery packet was received.
circuit-id remote-id: send-hostname-only	Setting this parameter to yes causes the system to send only the hostname and its length in the DHCP Option 82 suboption field. See “Including only the hostname identifier” on page 3-15. This parameter works in the same way for DHCP snooping and PPPoE snooping. For an example BRIDGE-GROUP configuration that enables the parameter, see “Sample bridge-group configuration for sending only a hostname” on page 3-19.
circuit-id remote-id: if-ip	IP address used to identify the Stinger system to the BRAS, AAA, or PPPoE server. If both circuit and remote ID subprofiles are enabled, only one address is used. If both this field and vendor-option-string are empty, the Stinger uses the system address if that value has been defined. This setting and the vendor-option-string setting are mutually exclusive in the same subprofile.
circuit-id remote-id: vendor-option-string	Text string, up to 31 characters, used to identify the Stinger system to the BRAS, AAA, or PPPoE server. This setting and the if-ip setting are mutually exclusive.
circuit-id remote-id: version	Configurable version ID, set to 1 by default.
circuit-id remote-id: dhcp-allow-any-src-port	Currently, this option is supported only for IPoA and BIR connections. For details, see “Interoperation with DHCP servers that zero-delimit suboption fields” on page 4-76.

Sample PPPOE snooping configuration

Figure 3-4 shows a Stinger system bridging user PVCs onto a VLAN using transparent bridging. The CPE is operating in bridging mode. In this case, PPPoE snooping is enabled in the bridge group, so DSL line identification will be added to upstream PPPoE Discovery packets.

Figure 3-11. PPPoE snooping example



To configure this VLAN for PPPoE snooping, first verify that bridging is enabled on the physical interface. For example:

```
admin> get ethernet { 1 8 2 } bridging-enabled
[in ETHERNET/{ shelf-1 first-control-module 2}:bridging-enabled]
bridging-enabled = yes
```

Then, follow these steps:

- 1 Create a bridge group, and enable PPPoE snooping. You must also configure the lan-router interface and configure either the remote-id or the circuit-id subprofile, or both. In this example, the circuit-id subprofile is configured, and a text string is used to identify the Stinger system as intermediate agent.

```
admin> new bridge-group
admin> set enable = yes
admin> set bridging-group = 451
admin> set lan-router-interface physical shelf = shelf-1
admin> set lan-router-interface physical slot = first-control-module
admin> set lan-router-interface physical item = 2
admin> set lan-router-interface logical-item = 50
admin> set pppoe-snooping enable = yes
admin> set pppoe-snooping circuit-id enable = yes
admin> set pppoe-snooping circuit-id vendor-option-string = stinger-001
admin> write -f
```

- 2 Create a VLAN-ETHERNET profile.

```
admin> new vlan-ethernet
admin> set interface-address physical shelf = shelf-1
admin> set interface-address physical slot = first-control-module
admin> set interface-address physical item = 2
admin> set interface-address logical-item = 50
admin> set enabled = yes
```


```
admin> set bridging-options bridging-group = 451
admin> set bridging-options bridge = yes
admin> set bridging-options bridge-type = transparent-bridging
admin> write -f
```

- 3 Create CONNECTION profiles for bridged PVCs to the CPE devices. The profiles must specify the right bridge group number.

```
admin> new connection
admin> set station = pppoe-client-1
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options ip-routing-enabled = no
admin> set bridging-options bridging-group = 451
admin> set bridging-options bridge = yes
admin> set bridging-options bridge-type = transparent-bridging
admin> set atm-options vpi = 0
admin> set atm-options vci = 36
admin> set atm-options nailed-group = 51
admin> write -f
```

With this configuration, when the client generates a PPPoE session request, the system adds the following options to the request packet:

Figure 3-12. Contents of fields with sample PPPoE snooping configuration



When the BRAS replies with options, the Stinger system removes those options before forwarding the reply to the PPPoE client.

Dynamic anti-spoofing for DHCP and PPPoE

Dynamic anti-spoofing prevents IP spoofing and PPPoE session spoofing. This feature uses the information (IP address or session ID) acquired during the DHCP initialization or PPPoE negotiations. Dynamic anti-spoofing is supported on connections that are part of transparent (N:1 VLAN) bridging. If multiple users are connected to the same interface (port/VP/VC), the `mac-address-learning-limit` parameter may be set in conjunction with the `dynamic-anti-spoofing` parameter to limit the number of DHCP/PPPoE negotiations.



Note This feature is supported only on IP2100.

IP anti-spoofing through DHCP

Stinger ensures that a customer cannot use an IP address if it is not attributed through DHCP. This is necessary to prevent service theft. Stinger compares the IP address retrieved from the DHCP OFFER to the source IP address in the traffic coming from the DSL side. If they do not match, Stinger drops the packets.

To support this feature, Stinger uses the Relay Agent Information option 82 suboption 1 (circuit-id) for the connections that have the dynamic-anti-spoofing enabled. When the DHCP client receives this option in the Request message, Stinger displays an error message. This option also has to be echoed back to the Stinger by the DHCP server. Stinger displays an error message if this option is not present in the reply.



Note You need to enable the diagnostics to view these messages.

After you enable the dynamic anti-spoofing, Stinger drops all traffic from this connection.

PPPoE anti-spoofing

Stinger ensures that a customer cannot use a PPPoE session if the customer has not initialized it. This is necessary to prevent service theft. The PPP Session ID retrieved from the PPPoE OFFER is compared to the PPPoE session ID in the traffic coming from the DSL side. If they do not match, Stinger drops the packets.

Configuring dynamic anti-spoofing

You can configure dynamic anti-spoofing per connection. By default the feature is disabled.

Once you enable the dynamic anti-spoofing on a connection, Stinger blocks the traffic until the DHCP initialize or the PPPoE negotiate. You must enable the DHCP-snooping or PPPoE-snooping on the bridge-group profile because Stinger has to snoop these messages to determine the IP address or the session ID. For connections where dynamic-anti-spoofing is enabled and DHCP-snooping or PPPoE-snooping are not enabled, a log message will be displayed when the connection comes up.

Enabling dynamic anti-spoofing

You can enable dynamic anti-spoofing by setting the `dynamic-anti-spoofing` parameter in the `bridging-options` subprofile to YES.

```
admin> read connection cpe-1
admin> set dynamic-anti-spoofing = yes
admin> write -f
```

Anti-spoofing aging

Anti-spoofing aging is a mechanism to delete the anti-spoofing parameters from the system. Anti-spoofing happens when:

- DHCP lease time expires and the client does not renew it.
- PPPoE session terminates abnormally.

In Stinger, anti-spoofing parameters age when the DHCP lease time expires and the client does not renew it. No configuration is required for this to happen.

The PPPoE anti-spoofing parameters age when there is no activity on the link for a period of time. The inactivity of the link is associated with the MAC address. Stinger allows initiation of multiple sessions from the same MAC address, but it will not detect inactivity if one session terminates abnormally and the other session is still up. The aging happens only when there is no activity on the link. By default the `mac-entry-age-time` is ON, and set to 300 seconds. If no activity is detected on the link within this time, Stinger ages the anti-spoofing parameters. The value of the `mac-entry-age-time` should be greater than the PPP LCP ECHO Request/Reply timer configured on the BRAS. If the client server does not have the PPP LCP ECHO Request/Reply mechanism, then `pppoe-anti-spoof-age-time` field in the `bridge-group` profile can be set to a value, which will be the inactivity timer. By default this field is set to 0.

If mac aging is disabled, then an internal timer is set to check DHCP lease expiry or PPPoE inactivity. The period is 60 seconds.

If the client server has the PPP LCP ECHO Request/Reply mechanism turned on, then the inactivity timer is same as the `mac-entry-age-time` in the `bridge-group`.

Setting anti-spoofing age time

Use the `pppoe-anti-spoof-age-time` parameter in the `bridge-group` profile to set the anti-spoofing aging time.

```
admin> read bridge-group 451
admin> set pppoe-anti-spoof-age-time = 300
admin> write -f
```

Spoofing parameters before and after aging

The following examples show how the anti-spoofing parameters look like for a DHCP negotiated IP call before and after aging.

Before Aging:

```
admin> info np spoof 35
dynamic-anti-spoofing :TRUE
FPP Rules for IP Spoof Entry:
Pattern IP-Address/Netmask
00000023 00000000      0.0.0.0/32
00000023 1E01FEFE     30.1.254.254/32  <== assigned source IP address
FPP Rules for PPPoE Spoof Entry:
Pattern      SessionId
00000023     0
```

After Aging:

```
admin> info np spoof 35
dynamic-anti-spoofing :TRUE
FPP Rules for IP Spoof Entry:
Pattern      IP-Address/Netmask
00000023 00000000      0.0.0.0/32
FPP Rules for PPPoE Spoof Entry:
Pattern      SessionId
00000023     0
```

The following examples show how the anti-spoofing parameters look like for a PPPoE type of call before and after aging.

Before Aging:

```
admin> info np spoof 35
dynamic-anti-spoofing :TRUE
FPP Rules for IP Spoof Entry:
Pattern                IP-Address/Netmask
00000023 00000000      0.0.0.0/32
FPP Rules for PPPoE Spoof Entry:
Pattern      SessionId
00000023      0
00230004      4  <== session id 4 is obtained from the server
```

After Aging:

```
admin> info np spoof 35
dynamic-anti-spoofing :TRUE
FPP Rules for IP Spoof Entry:
Pattern                IP-Address/Netmask
00000023 00000000      0.0.0.0/32
FPP Rules for PPPoE Spoof Entry:
Pattern      SessionId
00000023      0
```

Configuring stacked VLANs

VLAN stacking is a method of encapsulating one VLAN within another VLAN. It allows a carrier to partition the network among several network service providers (NSPs), while allowing each NSP to utilize VLANs to their full extent.

Each NSP can be assigned one or more VLANs (“backbone VLANs” or “NSP VLANs”), and within each NSP VLAN, up to 4095 unique instances of 802.1Q VLAN IDs are available, with each ID representing an ATM PVC from a DSL subscriber.

Stinger IP DSLAM systems support VLAN stacking for both untagged Ethernet frames and VLAN tagged traffic received on subscriber interfaces. The two methods differ in terms of subscriber CONNECTION profiles.




Note For stacked VLAN connections, the Stinger IP DSLAM does not bridge frames received from one DSL connection to another, even when the connections are configured with the same **bridging-group** value. This applies even to broadcast and multicast frames. Broadcast frames received from WAN interfaces are tagged with both the NSP VLAN and user VLAN ID and bridged to the appropriate NSP VLAN. Broadcast frames received from Gigabit Ethernet interface with a single VLAN tag are not bridged to the WAN interfaces.

Bridging untagged frames to stacked VLANs

In the sample stacked VLAN setup shown in Figure 3-13, a layer-2 core is partitioned among three NSPs.

Figure 3-13. Stacked VLAN: Bridging untagged frames from DSL interfaces



When the system receives untagged frames from an ATM PVC on a DSL interface, it tags the frames with the user's VLAN ID. This VLAN tag represents the user's connection within the context of an NSP. The system then embeds the tag in another, second-tier VLAN ID, which represents the user's NSP. The Ethernet frame from the ATM PVC is then bridged over the Gigabit Ethernet interface.

When system receives VLAN stacked frames on the Gigabit Ethernet interface, the NSP and subscriber VLAN tags are used to determine the DSL interface on which to bridge the frames. The two VLAN tags are removed before bridging the packet onto the WAN interface.

Overview of VLAN stacking settings for untagged frames

Each NSP requires one or more VLAN configurations on the Gigabit Ethernet interface, and each DSL subscriber requires a CONNECTION profile with a VLAN ID that is unique within the context of the user's destination NSP VLAN. Following are the parameters, shown with default settings, for VLAN stacking of untagged frames:

```
[in VLAN-ETHERNET/{ { shelf-1 first-control-module 2 } 50}:bridging-options]
bridge-type = vlan-circuit

[in CONNECTION/"" :bridging-options]
bridge-type = vlan-circuit
vlan-stack-user-vlan-id = 0

[in ETHERNET/{ shelf-1 second-control-module 2 }]
vlan-stack-tag-type = 91:00
```

Parameter	Setting
bridge-type	Type of bridging. Valid values are transparent-bridging, no-bridging, vlan-circuit, and stacked-vlan. For VLAN stacking, the stacked-vlan setting is required.

Parameter	Setting
vlan-stack-user-vlan-id	Subscriber's 802.1Q VLAN ID (from 0 to 4095) to be used in stacked-VLAN frames for incoming traffic that contains untagged Ethernet frames. The value must be unique within the NSP VLAN. With the default zero value, VLAN stacking is disabled for the connection. When set to a nonzero value, VLAN stacking is enabled and the specified value is added to the frames as a VLAN tag that represents the user's connection within the context of an NSP. This parameter is applicable only in CONNECTION profiles. Although it also appears in both ETHERNET and VLAN-ETHERNET profiles, the field is not applicable to those contexts.
vlan-stack-tag-type	Two-byte hexadecimal value to be inserted in the EtherType field for stacked-VLAN frames. All stacked VLAN frames being transmitted/received on the Ethernet port will use this value in their EtherType field. The default value is 0x9100. Because VLAN stacking is not yet a standardized technology, an EtherType value has not been standardized to represent stacked-VLAN frames, unlike the value 0x8100 used for IEEE 802.1Q VLAN frames, for example. So for the sake of interoperability, this value is configurable. If the layer-2 core network is using a specific EtherType value for stacked VLAN frames, you must set the <code>vlan-stack-tag-type</code> parameter to that value.

Sample configuration bridging untagged frames

The following example configures an NSP VLAN for “NSP 1” in Figure 3-13 (page 3-28), and two bridged PVCs to be directed to the NSP VLAN. In this example, the NSP VLAN is assigned VLAN number 50, and the user VLAN IDs are 471 and 473 within the NSP VLAN.

The following sets of sample commands configure a VLAN with multiple user PVCs.

- 1 Enable bridging on the Gigabit Ethernet port, as described in “Enabling layer 2 bridging for VLAN operations” on page 2-3. (If the layer-2 core network is using a specific EtherType value for stacked VLAN frames, you must also set the `vlan-stack-tag-type` parameter to that value.)
- 2 Create a VLAN-ETHERNET profile for the NSP VLAN. In this example, the VLAN is assigned VLAN ID 50 and `bridging-group 9`.

```
admin> new vlan-ethernet { { 1 8 2 } 50 }
admin> set enabled = yes
admin> set bridging-options bridging-group = 9
admin> set bridging-options bridge = yes
```

```
admin> set bridging-options bridge-type = stacked-vlan
admin> write -f
```


3 Create CONNECTION profiles for users of the NSP VLAN. These profiles must use the same bridging-group number as the NSP VLAN, and must specify a nonzero vlan-stack-user-vlan-id value that is unique within the NSP VLAN. In this example, the user VLAN IDs are 471 and 473.

```
admin> new connection vlan-user-1
admin> set encapsulation-protocol = atm
admin> set ip-options ip-routing-enabled = no
admin> set bridging-options bridging-group = 9
admin> set bridging-options bridge = yes
admin> set bridging-options bridge-type = stacked-vlan
admin> set bridging-options vlan-stack-user-vlan-id = 471
admin> set atm-options vci = 60
admin> set atm-options nailed-group = 125
admin> write -f
admin> new connection
admin> set station = vlan-user-2
admin> set encapsulation-protocol = atm
admin> set ip-options ip-routing-enabled = no
admin> set bridging-options bridging-group = 9
admin> set bridging-options bridge = yes
admin> set bridging-options bridge-type = stacked-vlan
admin> set bridging-options vlan-stack-user-vlan-id = 473
admin> set atm-options vci = 60
admin> set atm-options nailed-group = 127
admin> write -f
```

Bridging enterprise VLAN tagged frames to stacked VLANs

Stinger systems also support bridging of VLAN-tagged traffic received from the DSL side to stacked VLANs. The sample VLAN setup in Figure 3-14 (page 3-31) shows a Stinger IP DSLAM with three stacked VLAN configurations on the Gigabit Ethernet interface.

Figure 3-14. Stacked VLAN: Bridging enterprise VLAN-tagged frames



The CellPipe® units are operating in LLC bridged mode, and interacting with an enterprise VLAN-aware Ethernet switch in which enterprise VLANs are defined. Each enterprise IP subnet maps to a unique VLAN ID within the context of an individual DSL connection, and the Stinger IP DSLAM bridges the data between DSL interfaces and the Gigabit Ethernet backbone.

In the Stinger IP DSLAM, the basic VLAN stacking configuration on the Ethernet side does not change, but another layer of complexity is required to accommodate tagged frames from the DSL side.

Instead of adding a `vlan-stack-user-vlan-id` tag (as it does for untagged frames), when VLAN tagged frames are received from a DSL interface, the system must change the incoming VLAN tag to a VLAN user ID (and optionally, a new priority value) that represents the user's connection within the context of an NSP. The system then embeds the new (modified) VLAN information in another, second-tier VLAN tag, which represents the user's NSP. The Ethernet frame is then bridged over the Gigabit Ethernet interface.

The `vlan-stack-tag-type` parameter applies exactly as described for untagged frames in "Overview of VLAN stacking settings for untagged frames" on page 3-28.

Overview of VLAN stacking settings for tagged frames

Each NSP requires one or more VLAN configurations on the Gigabit Ethernet interface. To enable the system to modify VLAN tags in inbound tagged frames, the DSL CONNECTION profile must specify a FLOW-SERVICES profile that defines the mapping between enterprise priorities and VLAN IDs in inbound traffic and NSP priorities and VLAN user IDs for outbound traffic on the Gigabit Ethernet port. The profile can also be configured to influence switching or forwarding decisions on the basis of traffic flows. Each flow of traffic through an interface can be treated differently based on a set of flow classification rules.

Following are parameters, shown with default settings, for VLAN stacking of tagged frames:

```
[in FLOW-SERVICES/""]  
name* = ""  
service-type = none  
flow-list = [ { { 0 } { 0 } } { { 0 } { 0 } } { { 0 } { 0 } } { { 0 } { 0 } } ]+  
[in FLOW-SERVICES/"":flow-list]  
flow-list[1] = { { 0 0 00:00 } { 0 } }  
...  
flow-list[32] = { { 0 0 00:00 } { 0 } }  
[in FLOW-SERVICES/"":flow-list[1]]  
layer2-classifier = { 0 0 00:00 }  
service-options = { 0 }  
[in FLOW-SERVICES/"":flow-list[1]:layer2-classifier]  
vlan-id = 0  
ethernet-priority = 0  
ethernet-type = 00:00  
[in FLOW-SERVICES/"":flow-list[1]:service-options]  
vlan-id = 0  
[in CONNECTION/"":session-options]  
flow-services = ""
```

Parameter	Setting
name	Name of this FLOW-SERVICES profile, up to 31 characters. The name is used to identify the profile and apply it to a CONNECTION profile.
service-type	Type of service for the traffic flow of the connection. none With the default value of none, the profile has no effect on traffic flows if applied to a CONNECTION profile. vlan-stack-user-tag Enable the system to map the VLAN tag in the incoming traffic flow to a different VLAN user ID for stacked VLANs.
flow-list	An array of 32 indexed subprofiles used to configure a set of flow classification rules and corresponding service parameters applicable to those flows. The settings in the subprofiles are interpreted on the basis of the service-type setting. In a profile for tagged frames, the flow-list subprofiles must be configured in order (without a gap in the sequence of subprofile index numbers) beginning with flow-list 1, followed by flow-list 2, and so forth until all required flow-list subprofiles have been specified.

Parameter	Setting						
layer2-classifier	A flow-list subprofile for defining classification rules to identify a specific layer 2 traffic flow.						
	<table> <tr> <td>vlan-id</td><td>Enterprise VLAN ID in the VLAN-tagged frames of the inbound traffic flow (a number from 0 to 4095). This is the “source” VLAN ID that will be mapped to another ID specified in the service-options subprofile.</td></tr> <tr> <td>ethernet-priority</td><td>802.1Q/P priority value (p-bit value) of the frame’s VLAN tag. A number from 0 to 7.</td></tr> <tr> <td>ethernet-type</td><td>Ethernet type. A two-byte hexadecimal number representing the Ethernet protocol.</td></tr> </table>	vlan-id	Enterprise VLAN ID in the VLAN-tagged frames of the inbound traffic flow (a number from 0 to 4095). This is the “source” VLAN ID that will be mapped to another ID specified in the service-options subprofile.	ethernet-priority	802.1Q/P priority value (p-bit value) of the frame’s VLAN tag. A number from 0 to 7.	ethernet-type	Ethernet type. A two-byte hexadecimal number representing the Ethernet protocol.
vlan-id	Enterprise VLAN ID in the VLAN-tagged frames of the inbound traffic flow (a number from 0 to 4095). This is the “source” VLAN ID that will be mapped to another ID specified in the service-options subprofile.						
ethernet-priority	802.1Q/P priority value (p-bit value) of the frame’s VLAN tag. A number from 0 to 7.						
ethernet-type	Ethernet type. A two-byte hexadecimal number representing the Ethernet protocol.						
service-options	A flow-list subprofile for defining service-specific parameters applicable to the layer 2 traffic flow.						
	<table> <tr> <td>vlan-id</td><td>User VLAN ID (from 1 to 4095) to be inserted in stacked-VLAN frames. This is the “destination” VLAN ID that represents the user’s connection within the context of an NSP, and must be unique within the NSP VLAN.</td></tr> </table>	vlan-id	User VLAN ID (from 1 to 4095) to be inserted in stacked-VLAN frames. This is the “destination” VLAN ID that represents the user’s connection within the context of an NSP, and must be unique within the NSP VLAN.				
vlan-id	User VLAN ID (from 1 to 4095) to be inserted in stacked-VLAN frames. This is the “destination” VLAN ID that represents the user’s connection within the context of an NSP, and must be unique within the NSP VLAN.						
flow-services	Name of a FLOW-SERVICES profile that defines the proper VLAN tag mapping for this connection. Ethernet-priority and Ethernet-type are not considered for mapping to the User VLAN ID.						



Note In some cases, there may be a requirement to preserve the enterprise VLAN IDs received in incoming traffic and simply prefix them with the NSP VLAN ID. That can be done, as long as the enterprise VLAN IDs are unique within the context of the NSP VLAN, by specifying the enterprise VLAN ID in both the `layer2-classifier` and `service-options` subprofiles.

Sample configuration for mapping tagged frames

The following example configures an NSP VLAN for “NSP 2” in Figure 3-14 (page 3-31), and the bridged PVC from the unit labeled **CellPipe® 1**, to be directed to the NSP VLAN. The VLAN IDs in this example are shown in Table 3-6.

Table 3-6. Enterprise VLAN tagged frames mapped to a stacked VLAN

Enterprise VLAN ID	NSP VLAN ID	NSP VLAN user ID
1	478	1001
2	478	2002

Table 3-6. Enterprise VLAN tagged frames mapped to a stacked VLAN (Continued)

Enterprise VLAN ID	NSP VLAN ID	NSP VLAN user ID
3	478	3003

To configure the stacked VLAN and mapping between enterprise and NSP VLAN IDs, first verify that bridging is enabled on the physical interface. For example:

```
admin> get ethernet { 1 8 2 } bridging-enabled
[in ETHERNET/{ shelf-1 first-control-module 2}:bridging-enabled]
bridging-enabled = yes
```

If bridging is not enabled, enable it as described in “Enabling layer 2 bridging for VLAN operations” on page 2-3. Then, follow these steps:

- 1 Create a VLAN-ETHERNET profile for the NSP VLAN. In this example, the VLAN is assigned VLAN ID 478 and bridging-group 30.

```
admin> new vlan-ethernet { { 1 8 2 } 478 }
admin> set enabled = yes
admin> set bridging-options bridging-group = 30
admin> set bridging-options bridge = yes
admin> set bridging-options bridge-type = stacked-vlan
admin> write -f
```
- 2 Create a FLOW-SERVICES profile that maps VLAN tags 1, 2, and 3 on the DSL interface to user ID values that are unique within the NSP VLAN. In this example, the user VLAN IDs are 1001, 2002, and 3003.

```
admin> new flow-services
admin> set name = enterprise1-vlans
admin> set service-type = vlan-stack-user-tag
admin> set flow-list 1 layer2-classifier vlan-id = 1
admin> set flow-list 1 service-options vlan-id = 1001
admin> set flow-list 2 layer2-classifier vlan-id = 2
admin> set flow-list 2 service-options vlan-id = 2002
admin> set flow-list 3 layer2-classifier vlan-id = 3
admin> set flow-list 3 service-options vlan-id = 3003
admin> write -f
```
- 3 Create a CONNECTION profile for the inbound CellPipe® connection. The profile must use the same bridging-group number as the NSP VLAN, and must specify the FLOW-SERVICES profile.

```
admin> new connection cellpipe-1
admin> set encapsulation-protocol = atm
admin> set ip-options ip-routing-enabled = no
admin> set bridging-options bridging-group = 30
admin> set bridging-options bridge = yes
admin> set bridging-options bridge-type = stacked-vlan
admin> set session-options flow-services = enterprise1-vlans
```


```
admin> set atm-options vci = 60
admin> set atm-options nailed-group = 125
admin> write -f
```

Configuring routed VLANs

A routed VLAN interface is the interface to which the router's IP address on the VLAN is attached. In a Stinger IP DSLAM system, a routed VLAN interface is always mapped to a virtual IP interface on the Gigabit Ethernet port. Packets received on a routed VLAN interface are routed based on the IP address, and packets are sent through the routed VLAN interface based on an IP routing decision.

The system can route traffic received from a VLAN to any routed, terminated user connection, such as IPoA (RFC 2684), bridged IP routing (BIR), DHCP requests and responses, PPPoA, and PPPoE.

Figure 3-15. Sample routed VLAN



To configure a routed VLAN, follow these steps:

- 1 Create terminated routed CONNECTION profiles for subscribers.
For details see Chapter 4, “IP Routing Configuration.”
- 2 Create a VLAN on the Gigabit Ethernet interface.
The VLAN-ETHERNET parameters for routed VLAN are the same as for VLAN bridging (N:1). See “Overview of VLAN-ETHERNET and CONNECTION settings” on page 3-3.
- 3 Create a virtual IP interface to the configured VLAN, specifying the VLAN ID.



Note You can enable multicast forwarding on the virtual IP interface, to allow it to handle both routed and multicast traffic. For details about that configuration, see “Configuring MBONE interfaces” on page 9-3.

Creating a virtual IP interface for a routed VLAN

Following are the parameters, shown with default settings for a virtual IP interface, for configuring a routed VLAN:

```
[in IP-INTERFACE/{ { shelf-1 first-control-module 2 } 1 }]
interface-address* = { { shelf-1 first-control-module 2 } 1 }
ip-address = 0.0.0.0/0
vrouter = ""
vlan-enabled = no
```


vlan-id = 0

Parameter	Setting
interface-address	Physical address of the Gigabit Ethernet interface in the system, followed by the virtual IP interface number, using the following format: <code>{ shelf-n slot-n port-n } item-n }</code> The <code>item-n</code> value is a nonzero number identifying the virtual IP interface.
ip-address	IP address of the virtual IP interface.
vrouter	Name of a virtual router. With the default null value, the global router is used. If you specify the name of a virtual router, the specified virtual router is used for all traffic routed to that VLAN. For background information about configuring virtual routers, see Chapter 7, “Virtual Router Configuration.”
vlan-enabled	Enables or disables IEEE 802.1Q VLAN tagging on the virtual IP interface. Set this parameter to yes for a routed VLAN interface.
vlan-id	VLAN ID for this virtual interface. This is the IEEE 802.1Q tag value to be added to the IP packets transmitted on the virtual interface. The valid range is from 0 to 4095, but for full compatibility with IEEE 802.1Q, Alcatel-Lucent recommends that you do not use the <code>vlan-id</code> values of 0, 1 or 4095.

Sample routed VLAN configuration

Figure 3-16 shows a basic representation of a routed VLAN configuration, which enables the system to handle a routed data stream with a VLAN tag, and forward the data stream to CPE routers based on the CPE IP address.

Figure 3-16. Sample routed VLAN



To configure the routed VLAN interface, follow these steps:

- 1 Create a VLAN-ETHERNET profile. This example uses VLAN ID 101.
admin> new vlan-ethernet { { 1 8 2 } 101}
admin> set enabled = yes
admin> set bridging-options bridge-type = no-bridging
admin> write -f
- 2 Create an IP-INTERFACE profile with a valid IP address and VLAN enabled. Specify the same VLAN ID used in the VLAN-ETHERNET profile (101 in this example).
admin> new ip-interface { { 1 8 2 } 1 }
admin> set ip-address = 13.13.13.13/24
admin> set vlan-enabled = yes
admin> set vlan-id = 101
admin> write -f
- 3 Configure the terminated PVCs. The system routes packet streams to the virtual IP interface associated with the VLAN, so the VLAN is used to carry routed traffic. For example, the following commands configure PPPoA, IPoA, and BIR connections.
admin> new connection pppoa-1
admin> set active = yes
admin> set encapsulation-protocol = ppp
admin> set ip-options remote-address = 3.3.3.5/29
admin> set ppp-options send-auth-mode = chap-ppp-auth
admin> set ppp-options bi-directional-auth = required
admin> set ppp-options send-password = sendpw
admin> set ppp-options recv-password = recvpw
admin> set telco-options call-type = off
admin> write -f
admin> new connection ipoa-1
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 2.2.2.1/28
admin> set atm-options vci = 37
admin> set atm-options nailed-group = 301
admin> write -f
admin> new connection bir-1
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 4.4.4.5/32
admin> set ip-options local-address = 4.4.4.1/32

```
admin> set atm-options atm1483type = aal5-llc
admin> set atm-options vci = 111
admin> set atm-options nailed-group = 155
admin> set bir-options enable = yes
admin> write -f
```

Applying an IP filter to a routed VLAN

You can specify the name of an IP filter in the VLAN-ETHERNET profile to filter traffic routed to that VLAN. For details about creating and applying IP filters, see Chapter 11, “Filter Configuration.”

In the following sample commands, an IP filter is configured to drop traffic from the CPE connection named `ipoa-1` in Figure 3-15 (page 3-35), with the remote address 2.2.2.1/28.

```
admin> new filter out-from-cpe
admin> set output-filters 1 valid-entry = yes
admin> set output-filters 1 Type = ip-filter
admin> set output-filters 1 ip-filter protocol = 1
admin> set output-filters 1 ip-filter source-address-mask = 255.255.255.240
admin> set output-filters 1 ip-filter source-address = 2.2.2.1
admin> set output-filters 2 valid-entry = yes
admin> set output-filters 2 forward = yes
admin> write -f
```

The next commands apply the filter to routed VLAN 101:

```
admin> read vlan-ethernet { { 1 8 2 } 101 }
admin> set filter-name = out-from-cpe
admin> write -f
```

Assigning a virtual router to a routed VLAN

Virtual routing enables high-density circuit termination with secure logical partitioning and multiple route tables. Each virtual router represents an isolated domain, and interfaces belonging to that domain cannot be routed outside the domain.

If you specify the name of a virtual router in the `ip-interface` profile associated with a routed VLAN, that virtual router is used for all traffic routed to that VLAN. For details about configuring virtual routers, see the Chapter 7, “Virtual Router Configuration.”

The following sample commands configure a virtual router with a system address of 13.13.100.1:

```
admin> new vrouter vlan2
admin> set vrouter-ip-addr = 13.13.100.1
admin> write -f
```

The following commands assign the virtual router to a routed VLAN interface:

```
admin> read ip-interface { { 1 8 2 } 101 }  
admin> set vrouter = vlan2  
admin> write -f
```


The following commands specify that the CPE connections named pppoa-1 and bir-1 belongs in the vlan2 virtual routing domain:

```
admin> read connection pppoa-1  
admin> set vrouter = vlan2  
admin> write -f  
admin> read connection bir-1  
admin> set vrouter = vlan2  
admin> write -f
```

VLAN bridging of IPoA traffic to an upstream BRAS

Instead of using a terminating PVC for routed IP over ATM (IPoA) traffic, you can bridge IPoA traffic on a unique VLAN across Ethernet to terminate on an upstream BRAS system. You can implement this configuration without reconfiguring CPE devices and with minimal protocol overhead impact to facilitate migration to an Ethernet transport network.

Figure 3-17. Terminating IPoA traffic on an upstream BRAS



Upstream packet processing

In the upstream direction, the system receives routed IPoA packets from the DSL side (the ATM network). Before the packets are VLAN-circuit bridged to a BRAS on the Ethernet network, the system removes the LLC encapsulation and constructs an Ethernet header and 802.1Q VLAN tag.

To construct the Ethernet header, the Stinger DSLAM learns the MAC address of the BRAS by sending ARP requests over the configured VLAN, using a configured BRAS IP address. When it has obtained the required MAC address, the system begins bridging the routed IPOA packets to BRAS over the configured VLAN. For details about ARP on a VLAN, see “VLAN-based ARP table” on page 3-48.

Downstream packet processing

In the downstream direction, the system uses the VLAN ID of the incoming packet to retrieve the bridged IPoA connection parameters, such as VPI-VCI values, before removing the Ethernet header and VLAN tag and constructing the LLC-encapsulated IP packet.



Note The Stinger system does not examine the destination IP address in the downstream packets it receives from the BRAS.

Because the Stinger system does not examine the destination IP address in the packets it receives from the BRAS, the BRAS administrator must ensure that all IP data packets sent over a configured VLAN are intended for the CPE address, to prevent unauthorized access or unintended transmissions to the CPE. To accomplish this, the BRAS administrator can define an IP filter on the BRAS that allows VLAN traffic only to the CPE IP address. VLAN traffic to the Stinger IP control module (CM) GigE IP address in the BRAS profile should not be allowed. That address is used only for ARP purposes.

Configuration overview

To terminate an IPoA PVC on an upstream BRAS, the following configurations are required:

- Traffic from the IPoA connection must be mapped to a unique VLAN ID on the Ethernet transport network (1:1 VLAN bridging).
- The destination BRAS must be associated with one of the IP CM GigE ports.
A BRAS profile is indexed by the logical address of an IP CM GigE port, in which the logical item number is the BRAS ID number. A BRAS profile specifies:
 - The BRAS IP address.
 - A local IP address on the same subnet as the BRAS IP address. This address is used in ARP packets as the source IP address. It is not used for any other purpose.
- The connection's IPoA packets must be bridged rather than routed.
You must set the `packet-type` parameter to `ipoa` in the CONNECTION profile `bridging-options` subprofile.
- The destination BRAS must be associated with the CPE connection.
You must set the `bras-id` value in the CONNECTION profile `bridging-options` subprofile.

Overview of the BRAS profile

To bridge IPoA packets, a configured BRAS profile is required. Following are the parameter settings, shown with default values, identifying the destination BRAS and associating it with one of the IP CM GigE ports:

```
[in BRAS/{ { any-shelf any-slot 0 } 0 }]
interface-address* = { { any-shelf any-slot 0 } 0 }
enabled = no
ip-address = 0.0.0.0/0
bras-ip-address = 0.0.0.0/0
```


Parameter	Setting
interface-address	A logical address of an IP CM GigE port, in which the logical-item number is the BRAS ID. This address uses the following format: <pre>{ { shelf slot port } bras-id }</pre> For example, the following commands create a BRAS profile for an IP CM GigE port with BRAS ID 100: <pre>new bras { { 1 8 2 } 100 }</pre> The BRAS ID will be referenced in all CONNECTION profiles that terminate on this BRAS.
enabled	Enables or disables the BRAS profile.
ip-address	An IP address on the same subnet as the BRAS IP address. This address is used only in the context of this profile for the construction of VLAN-based ARP packets. For example: <pre>set ip-address = 1.1.1.1/24</pre> This address must be in the same subnet as the <code>bras-ip-address</code> for successful ARP resolution over a VLAN.
bras-ip-address	The IP address of the remote BRAS. <pre>set bras-ip-address = 1.1.1.132/24</pre>



Note The `ip-address` and `bras-ip-address` values must be on the same subnet. This requirement is not enforced by the system due to design decisions, but it is necessary for the ARP resolution process to succeed over a VLAN.

If the BRAS IP address is modified on the BRAS, you must modify the `bras-ip-address` in the BRAS profile. When you write the modified BRAS profile, the system bounces all connections that terminate on that BRAS, to populate the new `bras-ip-address` into all the related CONNECTION profiles.

Overview of CONNECTION profile settings

The following new parameters, shown with default values, are required in the IPoA CONNECTION profile:

```
[in CONNECTION/"" :bridging-options]  
packet-type = none  
bras-id = 0
```

Parameter	Setting
packet-type	The type of packets to be bridged. The valid values are <code>none</code> (the default), <code>ipoa</code>, and <code>pppoa</code>. When set to <code>ipoa</code>, the <code>bridging-group</code> parameter must also be configured. The meaning of the <code>packet-type</code> setting depends on the <code>bridge-type</code> setting.

VLAN Configuration

VLAN bridging of IPoA traffic to an upstream BRAS


Parameter	Setting
	packet-type bridge-type Packets bridged
	none vlan-circuit Bridged IP or PPP packets over VLAN
	none stacked-vlan Bridged IP or PPP packets over stacked VLAN
	ipoa ¹ vlan-circuit Routed IPOA packets over VLAN
	ipoa ² stacked-vlan Routed IPOA packets over stacked VLAN. (Not currently supported.)
	pppoa ³ vlan-circuit Routed PPPoA packets over VLAN. (Not currently supported.)
	pppoa stacked-vlan Routed PPPoA packets over stacked VLAN. (Not currently supported.)
bras-id	The BRAS ID, which is the final component of the index of a configured BRAS profile. For example, for the following BRAS profile: BRAS/{ { shelf-1 first-control-module 2 } 100 } The BRAS ID is 100.

1. If packet-type is set to ipoa and a bridging-group is not configured, an error message is displayed.
2. If packet-type is set to ipoa and bridge-type is not set to vlan-circuit, an error message is displayed.
3. If packet-type is set to pppoa, an error message is displayed.

Configuring VLAN bridged IPoA

Figure 3-18 shows a Stinger DSLAM bridging IPoA traffic from subscriber PVCs to a unique VLAN interface (1:1 VLAN bridging), to terminate on an upstream BRAS:

Figure 3-18. Sample 1:1 VLAN circuits for bridging IPoA



In this sample configuration, the routed IPoA traffic from CPE 1 is bridged to the BRAS over VLAN 100, and the routed IPoA traffic from CPE 2 is bridged to the BRAS over VLAN 101. For background information about 1:1 VLAN bridging, see “Configuring 1:1 VLAN bridging” on page 3-2.

- 1 Create a BRAS profile associated with an IP CM GigE port. For example:

```
admin> new bras { { 1 8 2 } 1 }
admin> set ip-address = 10.12.1.7/24
admin> set bras-ip-address = 10.12.1.1/24
admin> set enabled = yes
admin> write -f
```

- 2 Create a CONNECTION profile for CPE 1.

```
admin> new connection pvc-100
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options ip-routing-enabled = no
admin> set ip-options remote-address = 11.1.1.1
admin> set bridging-options bridging-group = 500
admin> set bridging-options bridge = yes
admin> set bridging-options packet-type = ipoa
admin> set bridging-options bridge-type = vlan-circuit
admin> set bridging-options bras-id = 1
admin> set atm-options vpi = 8
admin> set atm-options vci = 100
admin> set atm-options nailed-group = 51
admin> write -f
```

- 3 Create a VLAN-ETHERNET profile with the same bridge group number.

```
admin> new vlan-ethernet { { 1 8 2 } 100 }
admin> set enabled = yes
admin> set bridging-options bridging-group = 500
admin> set bridging-options bridge = yes
admin> write -f
```

- 4 Create a CONNECTION profile for CPE 2.

```
admin> new connection pvc-101
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options ip-routing-enabled = no
admin> set ip-options remote-address = 12.1.1.1
admin> set bridging-options bridging-group = 501
admin> set bridging-options bridge = yes
admin> set bridging-options bridge-type = vlan-circuit
admin> set bridging-options packet-type = ipoa
admin> set bridging-options bras-id = 1
admin> set atm-options nailed-group = 52
```

VLAN Configuration

VLAN bridging of IPoA traffic to an upstream BRAS

```
admin> set atm-options vpi = 8
admin> set atm-options vci = 101
admin> write -f
```


5 Create a VLAN-ETHERNET profile with the same bridge group number.

```
admin> new vlan-ethernet { { 1 8 2 } 101 }
admin> set enabled = yes
admin> set bridging-options bridging-group = 501
admin> set bridging-options bridge = yes
admin> write -f
```

Configuring bridged IPoA to multiple BRAS with the same IP address

Figure 3-19 shows a Stinger DSLAM bridging IPoA traffic to two upstream BRAS with the same IP address in different VLANs. The BRAS and the IP CM GigE logical interface are on the same subnet.

Figure 3-19. Two BRAS with the same IP address in different VLANs



With this network configuration, only one BRAS profile is required. All bridged IPoA connections terminating on either BRAS use the same BRAS ID in the CONNECTION profile. Because the Stinger system maintains a VLAN-based ARP table, ARP requests are sent to the appropriate BRAS based on the combination of the BRAS IP address and VLAN ID.

If the IP address of one of the BRAS is changed, you create a new BRAS profile and assign the new BRAS ID to connections that terminate on that BRAS.

To configure the sample network shown in Figure 3-19, follow these steps:

- 1 Create a BRAS profile associated with the IP CM GigE port. For example:

```
admin> new bras { { 1 8 2 } 2 }
admin> set ip-address = 10.10.1.22/24
admin> set bras-ip-address = 10.10.1.1/24
admin> set enabled = yes
admin> write -f
```

- 2 Create a VLAN-ETHERNET profile for VLAN 100.

```
admin> new vlan-ethernet { { 1 8 2 } 100 }
admin> set enabled = yes
admin> set bridging-options bridging-group = 525
admin> set bridging-options bridge = yes
admin> write -f
```

- 3 Create a CONNECTION profile for CPE 100.

The CPE connects to BRAS 2 through VLAN 100. Specify BRAS ID 2 and the bridging group number of VLAN 100.

```
admin> new connection pvc-100
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options ip-routing-enabled = no
admin> set ip-options remote-address = 11.11.1.1
admin> set bridging-options bridging-group = 525
admin> set bridging-options bridge = yes
admin> set bridging-options packet-type = ipoa
admin> set bridging-options bridge-type = vlan-circuit
admin> set bridging-options bras-id = 2
admin> set atm-options vpi = 8
admin> set atm-options vci = 100
admin> set atm-options nailed-group = 57
admin> write -f
```

- 4 Create a VLAN-ETHERNET profile for VLAN 101.

```
admin> new vlan-ethernet { { 1 8 2 } 101 }
admin> set enabled = yes
admin> set bridging-options bridging-group = 550
admin> set bridging-options bridge = yes
admin> write -f
```

- 5 Create a CONNECTION profile for CPE 200.

The CPE connects to BRAS 3 through VLAN 101. Specify BRAS ID 2 and the bridging group number of VLAN 101.

```
admin> new connection pvc-101
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options ip-routing-enabled = no
admin> set ip-options remote-address = 12.11.1.1
admin> set bridging-options bridging-group = 550
admin> set bridging-options bridge = yes
admin> set bridging-options bridge-type = vlan-circuit
```

VLAN Configuration

VLAN bridging of IPoA traffic to an upstream BRAS

```
admin> set bridging-options packet-type = ipoa
admin> set bridging-options bras-id = 2
admin> set atm-options nailed-group = 58
admin> set atm-options vpi = 8
admin> set atm-options vci = 101
admin> write -f
```

Caveat when configuring an unnumbered interface for bridged IPoA


This section illustrates the problems that can occur in the VLAN proxy ARP process when a bridged IPoA connection uses an unnumbered interface with certain network configurations. An unnumbered interface specifies the CPE address (`remote-address`) as the null address in a CONNECTION profile.

For a bridged IPoA connection with an unnumbered interface, the Stinger system responds to all ARP requests received on the VLAN from the BRAS, irrespective of the destination IP address in the received ARP request packet. This can cause unexpected results in the VLAN proxy ARP process.

Sample VLAN configuration

For example, Figure 3-20 shows a network configuration that is not recommended with the use of unnumbered interfaces. On the GigE side of the Stinger, two remote devices are in VLAN 10, and the CPE has a bridged IPoA connection to the BRAS with an unnumbered interface. In this case, the two remote devices are a BRAS and a PC, but the same principle applies for two BRAS.

Figure 3-20. Not recommended with unnumbered interfaces



Sample profiles

The following sample profiles show the settings of the CONNECTION, VLAN-ETHERNET, and BRAS profiles:

```
admin> new connection
admin> set station = ipoa-wan
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options ip-routing-enabled = no
admin> set ip-options remote-address = 0.0.0.0
admin> set bridging-options bridging-group = 500
admin> set bridging-options bridge = yes
```

```
admin> set bridging-options bridge-type=vlan-circuit
admin> set bridging-options packet-type = ipoa
admin> set bridging-options bras-id = 6
admin> set atm-options nailed-group = 51
admin> set atm-options vpi = 8
admin> set atm-options vci = 100
admin> write
admin> new vlan { 1 8 2 } 10 }
admin> set enabled = yes
admin> set bridging-options bridging-group = 500
admin> set bridging-options bridge = yes
admin> set bridging-options bridge-type=vlan-circuit
admin> write -f
admin> new bras { { 1 8 2 } 6 }
admin> set enabled = yes
admin> set ip-address = 10.10.1.11/16
admin> set bras-ip-address = 10.10.1.7/24
admin> write -f
```

Possible proxy ARP problems with the sample configuration

With the configuration shown above, if the BRAS tries to send data to the PC (or to another BRAS at 10.10.1.5, which is also in VLAN 10), the following events occur:

- the BRAS sends an ARP request over VLAN 10 for 10.10.1.5.
- The device at that address responds.
- The BRAS updates its ARP table with the MAC address for 10.10.1.5. The BRAS can then transmit data to that address.

However, because the CPE connection has an unnumbered interface, the following events also occur:

- The Stinger also replies to the BRAS ARP request over VLAN 10 with proxy ARP replies for the CPE, returning the IP CM GigE port MAC address.
- When the BRAS receives the ARP replies from the Stinger system, it updates its ARP table with the new MAC address for 10.10.1.5 and begins transmitting data to the Stinger instead of the PC.

In addition, another problem can occur if the PC sends a VLAN ARP request packet to the Stinger. If the PC is in the same subnet as the IP CM GigE port, the Stinger learns the MAC address of the PC and updates its VLAN ARP cache.

When many devices reside in the same subnet as the IP CM GigE port, and they are all on the same or different VLANs for which bridged IPoA connections exist, the VLAN ARP table on the Stinger system will be updated. This can cause a performance impact by populating the cache with many unnecessary entries.

Conclusions about the sample configuration

In summary, before configuring an unnumbered interface for a bridged IPoA connection, you should verify that the network configuration will not cause unexpected results of the VLAN proxy ARP process.

VLAN-based ARP table

Because the IPoA bridging subsystem can support multiple BRAS with the same IP address on different VLANs, the ARP table must distinguish duplicate addresses by VLAN ID. When the Stinger generates an ARP request to a BRAS over a configured VLAN interface, the ARP packet contains the following fields:

Table 3-7. ARP request packet contents on a configured VLAN

Field	Description
Source MAC address	MAC address of the IP CM GigE port.
Source IP address	Local IP address specified in the BRAS profile, which must be on the same subnet as the BRAS destination address.
Destination MAC address	Broadcast MAC address (FF:FF:FF:FF:FF:FF)
Destination IP address	BRAS IP address specified in the BRAS profile.

How the VLAN ARP entries are updated dynamically

When the system sends ARP requests to a BRAS, the following actions occur:

- A VLAN ARP cache is created in PENDING state, which exists for a maximum of 5 seconds if the ARP reply is not received.
If no ARP reply is received over the VLAN while the VLAN ARP cache state is in PENDING state, the VLAN ARP cache is deleted. Subsequent IPOA packets again trigger ARP requests over VLAN, creating a new VLAN ARP cache in PENDING state.
- If an ARP reply is received over the VLAN while the VLAN ARP cache state is in PENDING state, the cache state is updated to DYNAMIC state, which exists for 10 minutes.
- When the VLAN ARP cache has been in the DYNAMIC state for 9 minutes, the ARP subsystem sends 6 ARP request packets at one-second intervals to the BRAS to refresh the VLAN ARP cache.
If no ARP reply is received, the VLAN ARP cache is deleted.
- If the BRAS sends an ARP reply, the VLAN ARP cache state timer is reset to the latest value and the VLAN ARP cache is maintained in DYNAMIC state for another 10 minutes.

Adding VLAN ARP entries statically

You can add VLAN ARP entries statically by using the `arptable -a` command, which is described in more detail in the next section. When you specify the new VLAN ID argument, the entry is added to the VLAN ARP table.

If you add a VLAN ARP table entry statically by using `arptable -a`, the connections that terminate on the BRAS must be bounced manually to populate the new MAC address to the CONNECTION profiles.

Note that with this method of updating the VLAN ARP table, the hostname argument is the `bras-ip-address` from the BRAS profile, and the `ip-address` field of the BRAS profile is not used, so the system does not enforce the requirement that the `bras-ip-address` and `ip-address` must be in the same subnet. However, if they are not in the same subnet, proxy ARP on the VLAN will fail.

PPPoA bridging in 1:1 and N:1 VLAN

To bridge PPPoA traffic to an upstream BRAS or PPPoE server, the system converts incoming PPPoA traffic to PPPoE and initiates a new PPPoE session with the upstream network element. This feature is supported both for transparent bridged (N:1 VLAN) connections and for 1:1 VLAN circuits.

PPPoA bridging is supported for connections on both LIM and trunk interfaces. This feature is compliant with the DSL Forum draft *WT101/DSLForum2004.343*.

Upstream network element configuration requirements


The upstream network element configuration must meet the following requirements:

- The server must support multiple PPPoE sessions from same MAC address.
- The server must send PADT under the following conditions:
 - Detection of PPPoE session termination
 - Inconsistent session ID (either an inactive session ID or a session ID used by another MAC address)
 - The server must negotiate MRU at 1492 as needed by PPPoE and not a bigger MRU as negotiated by a PPPoA user. If the BRAS supports mini-jumbo frames, the MTU can be negotiated to a higher value. Currently, the Stinger system supports a maximum frame size of 1550 bytes, so the MRU should be negotiated accordingly
- The server must time out sessions.
- For N:1 VLAN bridging of PPPoA, the server must use vendor options in the PPPoE packets to identify an end user. This is not required for 1:1 VLANs, because the downstream interface is identified by the VLAN ID.

PPPoA bridging in N:1 VLANs

In Figure 3-21, bridge group 2111 supports both PPPoA and PPPoE CPE connections.

Figure 3-21. Sample transparent-bridged PPPoA configuration



To enable the system to bridge both the PPPoE and PPPoA data streams to the upstream server, you must enable PPPoA bridging in the bridge group. Without PPPoA bridging, the system discards PPPoA packets on a transparent bridging connection.

The source MAC address in the bridged PPPoA packets is the Stinger system's MAC address. Once the PPPoE session to the upstream server is established and the session ID is available, the Stinger system starts forwarding PPPoA packets upstream after adding the Ethernet, VLAN and PPPoE headers and Ethernet FCS as the trailer.

The bridging is performed on a point-to-point basis, not on the basis of MAC address. This means that, even though transparent bridging is enabled, no flooding can occur on the connection.

Required settings to enable PPPoA bridging in a bridge group

Following are the required parameters, shown with default values, for enabling PPPoA bridging in a bridge group:

```
[in BRIDGE-GROUP/0]
lan-router-interface-address = { { 1 0 0 } 0 }
pppoa-bridging-enabled = no
```

Parameter	Setting
lan-router-interface-address	When PPPoA bridging is enabled, the LAN router interface must specify the VLAN interface used to reach the upstream server.
pppoa-bridging-enabled	Specifies whether PPPoA bridging is enabled. With the default setting of no, PPPoA bridging is disabled. If you set this parameter to yes, both PPPoE and PPPoA data streams will be bridged upstream from this bridge group. Changes to the pppoa-bridging-enabled field are applied dynamically at run-time.

Downstream CPE configuration considerations

The PPPoA client configuration must meet the following requirements:

- The CPE configuration must either enable LQM or set up the idle timer. This is required to enable the Stinger system to detect if the upstream server has become unavailable, in order to terminate the session.

In addition, while PPPoA is being bridged on the CPE connection, the following traffic characteristics do not apply to the PPPoA traffic:

- Filtering
- Classification
- IGMP snooping
- DSL port blocking
- DHCP snooping

If the same connection can be used to carry either PPPoE or PPPoA, the Stinger system processes each type of data stream appropriately for bridging upstream. Note that no traffic flows occur between two PPPoA bridging DSL connections, regardless of the setting of the `port-block-enabled` parameter.

Sample configuration enabling PPPoA bridging in a bridge group

This section shows how to enable PPPoA bridging in a bridge group (N:1 VLAN) and provides an example of the following steps:

- 1 Configure a BRIDGE-GROUP profile.
- 2 Configure the VLAN interface to the upstream server.
- 3 Configure a transparent-bridged connection for PPPoA bridging.
- 4 Verify that a PPPoA session from the remote CPE is established properly.

Before you begin

Before you begin, complete these steps:

- 1 Check that bridging is enabled on the Gigabit Ethernet interface. For example:

```
admin> read ethernet { 1 8 2 }  
admin> set bridging-enabled = yes  
admin> write -f
```
- 2 Activate a DSL subscriber port. For example:

```
admin> read sdsl { 1 4 1 }  
admin> set enabled = yes  
admin> write -f
```

Create a bridge group

Create a BRIDGE-GROUP. You must specify an Ethernet “router” interface to the upstream server, and you must enable PPPoA bridging.

```
admin> new bridge-group  
admin> set bridging-group = 2111
```

VLAN Configuration

PPPoA bridging in 1:1 and N:1 VLAN

```
admin> set enable = yes
admin> set lan-router-interface physical shelf = shelf-1
admin> set lan-router-interface physical slot = first-control-module
admin> set lan-router-interface physical item = 2
admin> set lan-router-interface logical-item = 70
admin> set pppoa-bridging-enabled = yes
admin> write -f
```

Create a VLAN-ETHERNET profile

Create a VLAN-ETHERNET profile and configure it for transparent-bridging.

```
admin> new vlan-ethernet
admin> set interface-address physical shelf = shelf-1
admin> set interface-address physical slot = first-control-module
admin> set interface-address physical item = 2
admin> set interface-address logical-item = 70
admin> set enabled = yes
admin> set bridging-options bridging-group = 2111
admin> set bridging-options bridge = yes
admin> set bridging-options bridge-type = transparent-bridging
admin> write -f
```

Create a CONNECTION profile

Create a transparent-bridged CONNECTION profile that will carry PPPoA.

```
admin> new connection cpe-1
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options ip-routing = no
admin> set bridging-options bridge = yes
admin> set bridging-options bridging-group = 2111
admin> set bridging-options bridge-type = transparent-bridging
admin> set atm-options vpi = 0
admin> set atm-options vci = 25
admin> set atm-options nailed-group = 51
admin> write -f
```

Verify that the session is established

After the connections are up and the CPE has initiated a PPPoA session, you can use the new system-level `pppoaBriSess` command to see if the session has been “translated” from PPPoA to PPPoE for bridging purposes. For example:

```
admin> pppoaBriSess -a
BridgeGrp  BridgeType  IfNum  State          HostName
```

```
-----
      2111  TRANSPARENT      35  Active      cpe-1
```


To see more information about the session, use the `-i` option with the interface number. For example:

```
admin> pppoaBriSess -i 35
Translated PPPoA session details:
-----
Interface      = 35
Host Name      = cpe-1
Bridge Group   = 2111
State          = Active
PPPoE Session ID = 8
PPPoE Server Mac = 00:00:e2:19:bd:16
```

PPPoA bridging in 1:1 VLANs

In Figure 3-22, the PVC side of a 1:1 VLAN supports a PPPoA client.

Figure 3-22. Sample PVC-to-VLAN bridged PPPoA configuration



To enable the system to bridge the PPPoA data stream to the upstream server, you must specify that the connection carries the PPPoA packet type.

Once the PPPoE session to the upstream server is established and the session ID is available, the Stinger system starts forwarding upstream traffic received over the PVC to the network after adding the Ethernet, VLAN and PPPoE headers and Ethernet FCS as the trailer. Downstream PPPoA traffic is distinguished based on the VLAN ID.

Required setting to enable PPPoA bridging in a 1:1 VLAN

Following is the required parameter, shown with its default value, for enabling PPPoA bridging for a 1:1 VLAN:

```
[in CONNECTION/"" :bridging-options]
packet-type = none
```

Parameter	Setting
packet-type	Type of incoming packets on the LIM side. To enable PPPoA bridging, set this parameter to <code>pppoa</code>. This setting instructs the system to bridge routed PPPoA packets. With the default <code>none</code> setting, the connection bridges PPPoE packets. The <code>ipoa</code> packet type enables the bridging of routed IPoA packets.

Downstream CPE configuration considerations

The PPPoA client configuration must meet the following requirements:

- The CPE configuration must either enable LQM or set up the idle timer. This is required to enable the Stinger system to detect if the upstream server has become unavailable, in order to terminate the session.
- Filters and classification do not apply to the connection

If the same connection can be used to carry either PPPoE or PPPoA, the Stinger system processes each type of data stream appropriately for bridging upstream.

For 1:1 VLAN circuits, the system bridges either translated PPPoA or PPPoE based on the `packet-type` setting in the CONNECTION profile. If `packet-type` is set to `pppoa`, the system bridges translated PPPoA traffic. If `packet-type` is set to `none`, the system bridges translated PPPoE traffic.

Sample configuration enabling PPPoA bridging in a 1:1 VLAN circuit

This section shows how to enable PPPoA bridging in a 1:1 VLAN and provides an example of the following steps:

- 1 Configure the VLAN interface to the upstream server.
- 2 Configure a PVC that will carry PPPoA traffic.
- 3 Verify that a PPPoA session from the remote CPE is established properly.

Before you begin

Before you begin, complete these steps:

- 1 Check that bridging is enabled on the Gigabit Ethernet interface. For example:

```
admin> read ethernet { 1 8 2 }  
admin> set bridging-enabled = yes  
admin> write -f
```
- 2 Activate a DSL subscriber port. For example:

```
admin> read al-dmt { 1 5 1 }  
admin> set enabled = yes  
admin> write -f
```

Create a VLAN-ETHERNET profile

Create a VLAN-ETHERNET profile with VLAN circuit bridging

```
admin> new vlan-ethernet  
admin> set interface-address physical shelf = shelf-1  
admin> set interface-address physical slot = first-control-module  
admin> set interface-address physical item = 2  
admin> set interface-address logical-item = 21  
admin> set enabled = yes  
admin> set bridging-options bridge = yes  
admin> set bridging-options bridging-group = 1
```

```
admin> set bridging-options bridge-type = vlan-circuit
admin> write -f
```

Create a CONNECTION profile

Create a PVC CONNECTION profile that will carry PPPoA.

```
admin> new connection cpe-2
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options ip-routing = no
admin> set bridging-options bridge = yes
admin> set bridging-options bridging-group = 1
admin> set bridging-options bridge-type = vlan-circuit
admin> set bridging-options packet-type = pppoa
admin> set atm-options vpi = 0
admin> set atm-options vci = 100
admin> set atm-options nailed-group = 61
admin> write -f
```

Verify that the session is established

After the connections are up and the CPE has initiated a PPPoA session, you can use the new system-level `pppoaBriSess` command to see if the session has been “translated” from PPPoA to PPPoE for bridging purposes. For example:

```
admin> pppoaBriSess -a
```

BridgeGrp	BridgeType	IfNum	State	HostName
1	VLAN-CIRCUIT	27	Active	cpe-2
2111	TRANSPARENT	35	Active	cpe-1

To see more information about the session, use the `-i` option with the interface number. For example:

```
admin> pppoaBriSess -i 27
Translated PPPoA session details:
-----
Interface      = 27
Host Name      = cpe-2
Bridge Group   = 1
Bridge Type    = VLAN-CIRCUIT
State          = Active
PPPoE Session ID = 12
PPPoE Server Mac = 00:00:e2:19:bd:16
```

Administrative tools for VLAN

Commands that provide administrative information about VLAN are available only in the debug environment. If you are managing the system remotely, some of this information is also available through the `ip2kstats` MIB. For details, see “VLAN-related diagnostics” on page A-13 and “SNMP MIB for GMAC and VLAN statistics” on page A-29.

When bridging IPoA connections to an upstream BRAS, you can use the `arptable` command to get information about the VLAN-based ARP table.

The VLAN-based ARP table does not depend on virtual routers, and is maintained differently than the VLAN-based ARP table. For the sake of consistency, the syntax of `arptable` command is the same for both ARP tables, but the virtual router ID in the VLAN-based ARP table is ignored.

Table 3-8 shows the VLAN-specific arguments for the `arptable` command:

Table 3-8. VLAN ID argument added to arptable command

Command Syntax	Description of change
<code>arptable -a hostname MAC_address [vlan_id]</code>	The optional <code>vlan_id</code> argument has been added. If a VLAN ID is specified, the entry is added to VLAN-based ARP table. If a VLAN ID is not specified, the entry is added to the router-based ARP table for the appropriate virtual router.
<code>arptable -d hostname [vlan_id]</code>	The optional <code>vlan_id</code> argument has been added. If a VLAN ID is specified, the entry is deleted from the VLAN-based ARP table. If a VLAN ID is not specified, the entry is deleted from the router-based ARP table for the appropriate virtual router.

The `arptable` command displays both the router-based ARP table and VLAN-based ARP table entries. In addition, the `arptable -c` command now shows the entries of both ARP tables, and the `arptable -f` command flushes both ARP tables. For example:

```
admin> arptable
IP Address      MAC Address      Type  IF    Retries/Pkts/RefCnt  Time Stamp
135.112.124.32  00:B0:D0:A1:91:82  DYN   0      0/0/4649             146372
135.112.124.1   00:0C:DB:3C:11:00  DYN   0      0/0/3272             146840
Total arp entries (global and all vrouters): 2

IP Address      Van Id  MAC Address      Type  IF    Retries/Pkts/RefCnt  Time Stamp
150.150.100.60  10      00:11:22:33:44:55  DYN   0      0/0/4000             144544
```



```
150.150.100.60  20      00:11:22:33:44:66  DYN  0      0/0/4050
144587
150.150.100.40  30      00:11:22:33:44:77  STA  0      0/0/4100
145001
Total arp entries (All VLANs): 3
admin> arptable -c
Total arp entries (global and all vrouters): 2
Total arp entries (All VLANs): 3
admin> arptable -f
```

IP Routing Configuration

4

Introduction to the IP router software	4-1
Configuring IP-INTERFACE profiles for Ethernet ports	4-6
Configuring IP-GLOBAL network features	4-9
Configuring IP-ROUTE profiles	4-40
Overview of routed subscriber connection features	4-42
Configuring IPoA subscriber connections	4-49
Configuring BIR subscriber connections	4-56
Configuring DHCP relay for IPoA and BIR connections	4-65
Configuring broadband RAS subscriber access	4-85
Administrative tools for IP routing	4-109

This chapter describes IP routing features that are typically configured on a Stinger IP DSLAM. Some parameters in IP-related profiles are not used by the IP DSLAM, or are not relevant to its primary applications. Those parameters are not described in this chapter, but are documented in the *Stinger Reference*.

Introduction to the IP router software

When you reset the system, an IP routing table is constructed that contains all the routes known to the system, including the following:

- Routes for the local Ethernet interfaces (configured IP-INTERFACE profiles)
- Routes for active WAN IP sessions
- Routes for inactive WAN IP sessions (configured CONNECTION profiles)
- Routes defined in IP-ROUTE profiles or RADIUS route profiles

If the Routing Information Protocol (RIP) is enabled on one or more interfaces, the system adds routes as it learns them from routing-update packets. In addition, the system is continuously updating its routing table by adding routes for links that become active and removing routes for inactive sessions. If a nailed connection goes down, the system removes the route from its routing table.

Routes and interfaces

An IP route specifies a destination address, a gateway to the network, and an interface that leads to the gateway. It can also specify metrics and other values associated with the route.

A route defined in a profile is a *static route*. A *dynamic route* is learned from RIP updates sent by other routers. Dynamic updates provide access to many more routes than those actually configured in the system, and are updated automatically as routes change. However, dynamic updates cause additional routing overhead, so they are disabled by default.

An *interface* is a point of ingress to or egress from the system. For example, a local interface is an Ethernet port and a WAN interface is a nailed or switched connection. An *IP interface* is the logical IP address that enables IP data to be sent and received.

Displaying the routing table

For details about the `netstat` command, see the *Stinger Reference*. The following command displays the system's routing table:

```
admin> netstat -r
```

Destination	Gateway	IF	Flg	Pref	Met	Use	Age
0.0.0.0/0	1.112.26.1	ie0	SGP	60	1	343	2274
127.0.0.0/8	-	bh0	CP	0	0	0	2274
127.0.0.1/32	-	local	CP	0	0	0	2274
127.0.0.2/32	-	rj0	CP	0	0	0	2274
1.112.0.0/16	-	ie0	C	0	0	6497	2274
1.112.26.146/32	-	local	CP	0	0	3635	2274
224.0.0.0/4	-	mcast	CP	0	0	179	2274
224.0.0.1/32	-	local	CP	0	0	0	2274
224.0.0.2/32	-	local	CP	0	0	0	2274
224.0.0.9/32	-	local	CP	0	0	0	2274
255.255.255.255/32	-	ie0	CP	0	0	61	2274

Total Routes = 11 Hidden Routes = 0

For each route in the table, the **Destination** and **Gateway** fields show the destination address and the address of the next-hop router used to reach that destination. The zero destination address is the default route. If the system does not find a route for a packet's destination, it forwards the packet to the default route rather than dropping the packet. Note that the system uses the most specific route (having the longest prefix) that matches a given destination. Direct routes do not show a gateway address.

An asterisk (*) in the flags column indicates a hidden route, which is not included in routing updates sent by the system and is not used for forwarding packets. Hidden routes are used only for display purposes.

The **IF** field shows the name of the interface through which a packet addressed to the entry's destination will be sent. The route to the **mcast** interface name encapsulates the multicast forwarder for the entire class D address space. (For more information, see Chapter 9, "IP Multicast Configuration.")

Routes to the local unit display the **local** interface name. Packets to the 224.0.0.1 and 224.0.0.2 interfaces can be multicast and received like normal multicast packets, but upon receiving such a packet, the router does not forward it to another link layer device. Effectively, these packets have a maximum transmission unit (MTU) of 1.

Displaying the interface table

To display the interface table, use the `-i` option on the `netstat` command line:

```
admin> netstat -i
```

Name	MTU	Net/Dest	Address	Ipkts	Ierr	Opkts	Oerr
ie0	1500	1.112.0.0/16	1.112.26.146	5542	0	1636	0
ie1	1500	-	-	0	0	0	0
ie1-1	1500	-	-	0	0	0	0
lo0	1500	127.0.0.1/32	127.0.0.1	1629	0	1629	0
rj0	1500	127.0.0.2/32	127.0.0.2	0	0	0	0
bh0	1500	127.0.0.3/32	127.0.0.3	0	0	0	0
wanabe	1500	127.0.0.3/32	127.0.0.3	0	0	0	0
local	65535	127.0.0.1/32	127.0.0.1	1892	0	1892	0
mcast	65535	224.0.0.0/4	224.0.0.0	180	0	180	0
tunnel0	1500	1.112.0.0/16	1.112.26.146	0	0	0	0
vr0_main	1500	1.112.26.146/32	1.112.26.146	0	0	0	0
sip0	65535	-	-	0	0	0	0

The entries named `ie0`, `ie1`, or `ie1-N` represent Ethernet interfaces, where *N* represents the logical-item number of the interface. When the logical-item number is zero (the physical interface), it does not appear in the interface name. The other names in the interface table have the following significance:

- The `lo0` (loopback) interface is the local loopback.
- The `rj0` (reject) and `bh0` (blackhole) interfaces are used for summary routes, which are created by the system for pool summaries. The `rj0` interface is also used for summary routes created by the system for WAN virtual interfaces on bridged IP routing connections, when the virtual interface is assigned a prefix length of less than 32.
- The `wanabe` interface is an inactive RADIUS dial-out profile.
- The `local` interface is the local machine.
- The `mcast` interface is the multicast interface, which represents the multicast forwarder for the entire class D address space. For details, see Chapter 9, “IP Multicast Configuration.”
- The `tunnel` interface is a single pseudo-interface that is used only when the system is terminating tunnels. (The number terminating the tunnel interface name is an internal number that can change from one software version to the next.)
- The `vr0_main` interface represents the router itself.
- The `sip0` interface is the soft IP interface. For details, see “Defining a soft interface for increased accessibility” on page 4-8.
- The numbered WAN (`wanN`) interfaces are WAN connections, which are entered in the interface table as they become active.

IP control module performance statistics

The IP control module controller collects statistics on the number of packets and octets transmitted and received on each LIM interface. These counters are represented in the output of the `netstat -i` and `ifstat` commands, and are accessible to an external management utility.

The per-interface statistics for connections terminated on the IP control module are cleared on the LIM when they are displayed on the controller by using the `netstat -i` or `ifstat` command.

For example, in the following output, the `Ipkts`, `Ierr`, `Opkts`, and `Oerr` statistic counters display the sum of the packets in transit as seen by the IP control module added to the current packets counted by the TAOS interface manager.

```
admin> netstat -i
```

Name	MTU	Net/Dest	Address	Ipkts	Ierr	Opkts	Oerr
ie0	1500	10.1.26.0/24	10.1.26.1	605504	0	0	0
ie1	1500	15.1.1.0/24	15.1.1.1	0	0	0	0
lo0	1500	127.0.0.1/32	127.0.0.1	8	0	8	0
rj0	1500	127.0.0.2/32	127.0.0.2	0	0	0	0
bh0	1500	127.0.0.3/32	127.0.0.3	0	0	0	0
wanabe	1500	127.0.0.3/32	127.0.0.3	0	0	0	0
local	65535	127.0.0.1/32	127.0.0.1	58935	0	58935	0
mcast	65535	224.0.0.0/4	224.0.0.0	0	0	0	0
tunnel0	1500	10.1.26.0/24	10.1.26.1	0	0	0	0
vr0_main	1500	10.1.26.1/32	10.1.26.1	0	0	0	0
sip0	65535	-	-	0	0	0	0
wan11	1524	200.200.200.254	2.2.2.1	7	0	10	0

Similarly, in the following output, the `in_oct`, `in_errs`, `out_octet`, and `out_err` statistic counters display the sum of the packets in transit as seen by the IP control module added to the current packets counted by the TAOS interface manager.

```
admin> ifstat 1
in_oct 0 in_errs 0 out_octet 0 out_err 0
```

The MIB II interface stat counters also now display the correct values when viewed from an external management utility.

IP address syntax

The system uses dotted decimal format (not hexadecimal) for IP addresses. If no subnet mask is specified, the system uses a default mask based on the address class. For example, a class C address, such as 198.5.248.40, has 24 network bits, leaving 8 bits for the host portion of the address. If no subnet mask is specified for a class C address, the system uses the default mask of 24 bits.

A subnet address includes a prefix length, which specifies the number of network bits in the address. For example, the following address specifies a 29-bit subnet:

```
ip-address = 198.5.248.40/29
```

In this address, 29 bits of the address are used to specify the network. The three remaining bits are used to specify unique hosts on the subnet. With three bits used to specify hosts on a 29-bit subnet, eight different bit combinations are possible. Of those eight possible host addresses, two are reserved:

- 000 — Reserved for the network (base address)
- 111 — Reserved for the broadcast address of the subnet



Note Be careful with zero subnets (subnets with the same base address as a class A, B, or C network). Early implementations of TCP/IP did not allow them. For example, the subnet 192.32.8.0/30 was illegal because it had the same base address as the class C network 192.32.8.0/24, while the subnet 192.32.8.4/30 was legal. Modern implementations of TCP/IP support zero subnets, and the Stinger implementation of RIP treats these subnets the same as any other network. However, you must treat zero subnets consistently throughout your network. Otherwise, you will encounter routing problems.

Table 4-1 shows subnet masks and prefix lengths for a class C network number.

Table 4-1. Decimal subnet masks and corresponding prefix lengths

Subnet mask	Number of host addresses	Prefix length
255.255.255.0	254 hosts + 1 broadcast, 1 network base	/24
255.255.255.128	126 hosts + 1 broadcast, 1 network base	/25
255.255.255.192	62 hosts + 1 broadcast, 1 network base	/26
255.255.255.224	30 hosts + 1 broadcast, 1 network base	/27
255.255.255.240	14 hosts + 1 broadcast, 1 network base	/28
255.255.255.248	6 hosts + 1 broadcast, 1 network base	/29
255.255.255.252	2 hosts + 1 broadcast, 1 network base	/30
255.255.255.254	Invalid mask (no hosts)	/31
255.255.255.255	1 host—a host route	/32

The broadcast address of any subnet has the host portion of the IP address set to all ones. The network address (or base address) represents the network itself, because the host portion of the IP address is all zeros. For example, supposing the IP configuration assigns the following address to a remote router:

198.5.248.120/29

The Ethernet network attached to that router has the following address range:

198.5.248.120 – 198.5.248.127

A host route is a special-case IP address with a prefix length of /32. For example:

198.5.248.40/32

Host routes are to a single host, rather than to a router or subnet.

Configuring IP-INTERFACE profiles for Ethernet ports

The system creates an IP-INTERFACE profile for an Ethernet port when it first detects the presence of the port.

Table 4-2. Profiles for configuring logical IP interfaces for Ethernet ports

Interface #	Type of interface	System-generated configuration profiles
0	Soft interface	IP-INTERFACE/{ { any-shelf any-slot 0 } 0 }
1	IP2000 10/100 BASE-T (management)	IP-INTERFACE/{ { shelf-1 first-control-module 1 } 0 }
2	IP2000 Gigabit Ethernet	IP-INTERFACE/{ { shelf-1 first-control-module 2 } 0 }
0	Soft interface	IP-INTERFACE/{ { any-shelf any-slot 0 } 0 }
1	IP2100 10/100 BASE-T (management)	IP-INTERFACE/{ { shelf-1 first-control-module 1 } 0 }
2	IP2100 Gigabit Ethernet	IP-INTERFACE/{ { shelf-1 first-control-module 2 } 0 }
3	IP2100 Gigabit Ethernet	IP-INTERFACE/{ { shelf-1 first-control-module 3 } 0 }

This index consists of a physical address and a logical-item number in the following format:

{ { shelf-num slot-num item-num } logical-item-num }

The logical item addresses a specific logical interface. It is zero except when multiple (virtual) interfaces have been configured on the physical port. For more details, see “Defining a local virtual IP interface” on page 4-8.

Overview of typical local interface settings

For information about enabling IP multicast forwarding on the Gigabit Ethernet interface, see Chapter 9, “IP Multicast Configuration.”

Following are the parameters, shown with default settings, used to configure the Gigabit Ethernet port as an IP interface:

```
[in IP-INTERFACE/{ { shelf-1 first-control-module 2 } 0 }]
interface-address* = { { shelf-1 first-control-module 2 } 0 }
ip-address = 0.0.0.0/0
rip-mode = routing-off
rip2-use-multicast = yes
directed-broadcast-allowed = yes
vlan-enabled = no
vlan-id = 0
```

Parameter	Setting
interface-address	Address of the interface in the Stinger unit, or, if the item number is not zero, the virtual interface address.

Parameter	Setting
ip-address	IP address of the LAN interface. If the LAN IP address includes a subnet specification, you must create a static route to another LAN router to enable the system to reach local networks beyond its own subnets. For details, see “Configuring IP-ROUTE profiles” on page 4-40.
rip-mode	Enables or disables RIP updates on the interface. You can enable RIP to receive routing table updates, send them, or both. Running RIP-2 and RIP-v1 on the same network in such a way that the routers receive each other’s advertisements is <i>not recommended</i> .
rip2-use-multicast	Enables or disables use of the multicast address (224.0.0.9) rather than the broadcast address for RIP updates.
directed-broadcast-allowed	Enables or disables forwarding of directed broadcast traffic onto the interface and its network.
vlan-enabled vlan-id	Configure a routed VLAN interface. See “Configuring routed VLANs” on page 3-35.

Configuring a local IP interface

The following command assigns an IP address to the first Gigabit Ethernet port of the first control module (installed in slot 8):

```
admin> read ip-interface { { 1 8 2 } 0 }  
admin> set ip-address = 1.1.1.2/29  
admin> write -f
```

In this example configuration, the Gigabit Ethernet interface is connected to the 1.1.1 subnet. RIP is off by default, so to enable the interface to communicate with routers on other local subnets, either the system must have a static route configuration to another router in its own subnet, or the interface must enable RIP. For an example of configuring a static route, see “Configuring IP-ROUTE profiles” on page 4-40.

The following commands configure the interface to receive RIP-2 updates on the multicast address (the multicast address is the default):

```
admin> read ip-interface { { 1 8 2 } 0 }  
admin> set rip-mode = routing-recv-v2  
admin> write -f
```

You can verify that the system can transfer IP packets across the interface by pinging another host on the same network segment, as shown in the following example:

```
admin> ping 1.1.1.19  
PING 1.1.1.19: 56 Data bytes  
64 bytes from 1.1.1.19: icmp_seq=0 ttl=255 time=0 ms  
...
```

Defining a local virtual IP interface

You can configure up to 16 IP-INTERFACE profiles for each Ethernet port, with each profile specifying one IP address. However, if you have a link aggregation group (LAG) interface configured, the maximum number of virtual interfaces decreases by 1.

For details about using a virtual IP interface for a routed VLAN, see “Configuring routed VLANs” on page 3-35.

The system creates the default profile for each interface and assigns it the zero logical-item number. To configure another IP address on a LAN interface, create an IP-INTERFACE profile with a nonzero logical-item number in its interface address. For example:

```
admin> read ip-interface { { 1 8 2 } 1 }
admin> set ip-address = 1.1.1.1/29
admin> write -f
```

The logical-item numbers do not have to be consecutive, but they must each be unique.

Defining a soft interface for increased accessibility

You can configure a soft IP interface, which is an internal IP interface that is always active and reachable, as long as one of the system’s IP interfaces is up. The IP-INTERFACE profile with the zero index is reserved for the soft interface.



Note Do not specify the IP address of a physical LAN interface as the soft interface address.

The following commands set the soft interface IP address to 1.1.1.128:

```
admin> read ip-interface { { 0 0 0 } 0 }
admin> set ip-address = 1.1.1.128
admin> write -f
```

If RIP is enabled, the system advertises the soft interface address as a host route (with a prefix length of 32 bits) using the loopback interface. If RIP is not enabled, routers one hop away from the unit must have a static route to the soft interface address.

To verify that other hosts in your network have a route to the soft address, run `ping` or `tracert` from the other hosts. For example:

```
host1% ping 1.1.1.128
PING 1.1.1.128 (1.1.1.128): 56 Data bytes
64 bytes from 1.1.1.128: icmp_seq=0 ttl=255 time=0 ms
...
```

Disabling directed broadcasts to protect against denial-of-service

Denial-of-service attacks known as *smurf* attacks typically use ICMP Echo Request packets with a spoofed source address and packets directed to IP broadcast addresses. These attacks are intended to degrade network performance, possibly to the point that the network becomes unusable.

To prevent the IP router from being used as an intermediary in this type of denial-of-service attack launched from another network, you must disable the router from forwarding directed broadcasts it receives from another network. You must explicitly disable directed broadcasts on *all* IP interfaces in the system (including the management interface). In a system with redundant controllers, disable the feature on both controllers, so the unit is still protected following a switchover. The following commands configure the Gigabit Ethernet interface on the first controller:

```
admin> read ip-interface { { 1 8 2 } 0 }  
admin> set directed-broadcast-allowed = no  
admin> write -f
```

Configuring IP-GLOBAL network features

The IP router has many configuration settings that affect its operations. The settings that determine its routing policies include security, RIP options, IP route cache options, and other options. These settings are available only in the IP-GLOBAL profile. They have no counterpart in RADIUS.

Of the many settings in the IP-GLOBAL profile, some are intended for remote access service and are not directly relevant to IP DSLAM operations. However, there are many more options you can choose to configure in this profile. For details about all IP-GLOBAL parameters and subprofiles, see the *Stinger Reference*.

Setting a system address

The system IP address is the source address used for all packets generated by the system. It must be the real address of one of the unit's LAN IP interfaces, or the soft interface address (see "Defining a soft interface for increased accessibility" on page 4-8.) Following is the parameter for specifying a system address:

```
[in IP-GLOBAL]  
system-ip-addr = 0.0.0.0
```

With the default zero address, the Stinger unit uses the IP address assigned to the management IP address (configured in the IP-INTERFACE profile with the {1 f 1 } index), or if that address is not specified, the Gigabit Ethernet interface({ 1 f 2 }) as the source address for packets it generates. However, explicitly setting the system address simplifies access control. For example, most RADIUS servers keep a database of known remote access server (RAS) clients and their authentication keys. If you do not specify a system address, the RADIUS database must include a complete list of all the system's interface addresses. If you specify a system address, it is used for all RADIUS request packets.

Following is an example of setting the system-ip-addr parameter to the Ethernet interface address:

```
admin> get ip-interface { { 1 8 2 } 0 } ip-address  
ip-address = 2.2.2.2  
admin> read ip-global  
admin> set system-ip-addr = 2.2.2.2  
admin> write -f
```

Configuring DNS

Domain Name System (DNS) is a TCP/IP service for centralized management of address resolution. You enable DNS lookups by specifying a domain name and the IP addresses of one or more local servers.

Some sites maintain multiple DNS servers, each one dedicated to a particular client or location. In addition, some servers support a list feature that enables them to return multiple addresses for a hostname in response to a DNS query. For information about those DNS features, see the *Stinger Reference*.

Overview of typical DNS settings

Following are the parameters (shown with default settings) for configuring DNS to allow lookups:

```
[in IP-GLOBAL]
domain-name = ""
dns-primary-server = 0.0.0.0
dns-secondary-server = 0.0.0.0
sec-domain-name = ""
```

Parameter	Setting
domain-name	Primary domain name to use for DNS lookups. The system appends this domain name to hostnames when performing lookups.
dns-primary-server	Address of the primary local DNS server to use for lookups.
dns-secondary-server	Address of the secondary local DNS server to use for lookups. Used only if the primary server is not found.
sec-domain-name	Secondary domain name to use for DNS lookups if the hostname is not found in the primary domain.

Specifying domain names for lookups

The following commands specify a primary and secondary domain name for DNS lookups:

```
admin> read ip-global
admin> set domain-name = abc.com
admin> set sec-domain-name = eng.abc.com
admin> write -f
```

If a lookup fails with the first domain name, the router tries again with the secondary domain name.

Setting RIP options

The following parameters (shown with default settings) define how the system handles RIP updates:

```
[in IP-GLOBAL]
rip-policy = Poison-Rvrs
summarize-rip-routes = no
rip-trigger = yes
rip-pref = 100
dialout-poison = no
rip-queue-depth = 0
ignore-def-route = yes
suppress-host-routes = no
```

Parameter	Setting
rip-policy	Policy for sending update packets that include routes received on the same interface.
summarize-rip-routes	Enables or disables summarization of network information in RIP-v1 updates. RIPv1 does not advertise a netmask. This setting has no effect on RIP-2 updates. If enabled, the system “rounds up” the advertisement to the class of network being advertised. For example, for the route 10.10.10.0/24, RIPv1 would advertise 10.0.0.0. If <code>summarize-rip-routes</code> is not enabled, RIPv1 would advertise 10.10.10.0.
rip-trigger	Enables or disables RIP triggering. With a <code>yes</code> setting (the default), RIP updates include only changed routes.
rip-pref	Default preference for routes learned from RIP updates. When choosing the routes to put in the routing table, the unit first compares their preference values, preferring the lowest number. If the preference values are equal, the router compares the metric values, using the route with the lowest metric. Specify a number from 0 through 255. A value of 255 prevents the use of the route. The default preferences for different types of routes are 0 (zero) for connected routes, 10 for OSPF routes, 30 for routes learned from ICMP redirects, and 100 for routes learned from RIP and static routes.
dialout-poison	Enables or disables advertisement of dial-out routes when no trunks are available. Stinger units do not dial out, so leave this parameter at its default setting.
ignore-def-route	Enables or disables exclusion of advertised default routes from the routing table.
rip-queue-depth	Maximum number of RIP packets to be held for processing. Valid values are 0 to 1024. The default (0) means that the router will not drop any RIP packets, no matter how far behind it gets.
suppress-host-routes	Enables or disables suppression of host routes for interfaces with a prefix length of less than 32 bits.

RIP policy for propagating updates back to the originating subnet

You can specify a split-horizon or poison-reverse policy for outgoing update packets that include routes received on the same interface on which the update is sent. Split-horizon means that the router does not propagate routes back to the subnet from which they were received. Poison-reverse means that it propagates routes back to the subnet from which they were received, but with a metric of 16 (infinite metric).

The following set of commands specifies the split-horizon policy:

```
admin> read ip-global
admin> set rip-policy = split
admin> write -f
```

RIP triggering

RIP triggering enables the router to tag routes that have been updated in the routing table and send updates that include only the changed routes. The result is reduced processing overhead for both the TAOS router and its neighbors.

With the default value (yes), the router tags changes to its routing table and includes only the tagged routes in its next update. Changes occur when a call arrives or disconnects, RIP learns a route from another router, or the administrator modifies a route-related profile. The router broadcasts updates 5 to 8 seconds after the first change in the routing table is detected. The delay helps to prevent constant updates during peak traffic conditions.

If `rip-trigger` is set to `no`, the router sends full table updates every 20 to 40 seconds. To prevent RIP routers on a network from synchronizing and sending large updates in unison, the full table update is no longer broadcast at fixed 30-second intervals.

Limiting the size of UDP packet queues

When the router is very busy and receives a flood of UDP packets from SNMP requests or RIP updates, a backlog of packets waiting for processing can create enough delay in routing to cause sporadic problems with time-sensitive packets, such as LCP negotiation or frame relay management packets.

To prevent such problems, UDP processing runs at a lower priority than processing of routed packets. On a system busily routing packets, UDP processing might be delayed, and a backlog of UDP packets builds up. The `rip-queue-depth` parameter in the IP-GLOBAL profile and the `queue-depth` parameter in the SNMP profile specify the maximum size of this backlog.

When you set one of these parameters to specify a queue depth, the system is more likely to drop UDP packets when it is busy routing packets. However, time-sensitive routed packets are less likely to be delayed and system memory is used more efficiently.

In following sample commands sets both queue depths to 50. Fifty of each type of packet is held for processing, and if additional packets of either type are received when the queue is full, they are dropped.

```
admin> read ip-global
admin> set rip-queue-depth = 50
admin> write -f
```

```
admin> read snmp
admin> set queue-depth = 50
admin> write -f
```

The `netstat` command output shows the queue depth of various UDP ports, and the total packets received and total packets dropped on each port. The total packets received count includes dropped packets. In the following example, the SNMP queue depth was set to 32:

```
admin> netstat udp
udp:
Socket  Local Port  InQLen  InQMax  InQDrops  Total Rx
0       1023       0       1       0         0
1       route     0       50      0         509
2       echo      0       32      0         0
3       ntp       0       32      0         0
4       1022      0       128     0         0
5       SNMP     32      32     5837     20849
```

Ignoring default routes when updating the routing table

Alcatel-Lucent recommends enabling the `ignore-def-route` parameter to prevent routing updates from modifying the default route in the routing table. The following set of commands protects the default route from RIP updates:

```
admin> read ip-global
admin> set ignore-def-route = yes
admin> write -f
```

Suppressing host-route advertisements

If you set the `suppress-host-routes` parameter to `yes`, host routes (/32) are suppressed according to the following rules:

- If a CONNECTION profile includes a prefix length of less than 32 bits in a `remote-address` setting, host routes for the interface are suppressed while the session is being negotiated, and after the session is established, only the summary route is advertised for the interface. For examples, see “Overview of summary routes” on page 4-27 and “How routes are created for WAN virtual interfaces” on page 4-65.
- If a CONNECTION profile includes a subnet mask of /32 in the `remote-address` setting, host routes for the interface are not suppressed. (Pool addresses also have a 32-bit mask, so they are not suppressed.)

The following set of commands configures the router to suppress host routes:

```
admin> read ip-global
admin> set suppress-host-routes = yes
admin> write -f
```

Subscriber profile sharing

All Stinger LIMs support profile sharing for incoming subscriber connections. With shared profiles, each session is a separate connection that shares the same name and password. If an IP address is assigned, it must be assigned dynamically so that each session has a unique address.



Note Because the password is shared by multiple callers, profile sharing is unsuitable for high-security networks.

You enable profile sharing on one of two possible levels: globally in the IP-GLOBAL profile or per-connection in a CONNECTION profile. Once profile sharing is enabled globally, you cannot disable it for an individual connection. You enable profile sharing through the use of the following parameters, shown with their default settings:

```
[in IP-GLOBAL]
shared-prof = no
[in CONNECTION/"" ]
shared-prof = no
max-shared-users = 0
```

With the default settings, profile sharing is prohibited.

With the following settings, profile sharing is enabled systemwide, and the **shared-1** connection profile limits the number of concurrent sessions to 10:

```
admin> read ip-global
admin> set shared-prof = yes
admin> write -f
admin> read connection shared-1
admin> set shared-prof = no
admin> set max-shared-users = 10
admin> write -f
```

With the following settings, profile sharing is disabled systemwide, but the **shared-1** CONNECTION profile enables profile sharing and limits the number of concurrent sessions to 10:

```
admin> read ip-global
admin> set shared-prof = no
admin> write -f
admin> read connection shared-1
admin> set shared-prof = yes
admin> set max-shared-users = 10
admin> write -f
```


Equal cost multipath routing

The system supports load balancing of up to four equal cost multipath (ECMP) routes to a given destination. You can configure ECMP routes statically, or the system can learn ECMP routes through RIP updates. ECMP routes must specify the same destination address and different gateway addresses. Currently, costs and metrics are not considered in determining that a route is ECMP.

The gateway address of a route specifies the IP address of a next-hop router used to reach the destination. In the upstream direction, the gateway is a router connected to the same Ethernet segment as the Stinger system. In the downstream direction, the gateway is a router that is one hop away on a WAN link.

- Upstream, ECMP routing provides redundant paths to a destination.
- Downstream, ECMP routing provides an alternative to multilink PPP by enabling you to bind up to four links to one subscriber.

ECMP routing is supported through all IP interfaces or subinterfaces of the system, including routed VLANs, IP-terminated connections, multiple Gigabit Ethernet ports (if available) and so forth. If more than one Gigabit Ethernet port is supported in the system, ECMP routing is supported through all IP interfaces or subinterfaces of all the ports.

Packet-based ECMP

With packet-based load balancing on ECMP routes, the system randomly distributes the packets among the redundant paths. Test have shown that this method results in a fairly even load distribution over a period of time.

Flow-based ECMP

With flow-based load balancing on ECMP routes, you define a flow as a class of traffic based on one or more of the following fields in packet headers:

- destination IP address
- source IP address
- IP protocol type

Packets belonging to the same unicast flow always take the same path. Load balancing is on a best-effort basis, and becomes fairly uniform in distribution only when the number of flows is relatively large.

With flow-based load balancing, if you do not specify a packet header field to define a flow, the system selects randomly on which gateway to forward an ECMP packet stream, and thereafter all the packets flow through that interface.

In this event, you can force the ECMP subsystem to reallocate the interface for a packet stream by setting `ecmp-type` to `load-balanced` and writing the profile, and then setting `ecmp-type` back to `flow-based` and writing the profile.

Caveat about changing the type of load balancing

If you change `ecmp-type` value from `load-balanced` to `flow-based` (or vice versa) after ECMP routes have been added to the system, the existing ECMP routes continue to use the type of load-balancing that was previously in effect. When the two types of ECMP load balancing are in use on different sets of ECMP routes, it becomes difficult to track the behavior of the routing mechanism.

Limitations of ECMP load balancing

The following limitations apply with this software version:

- The system does not currently look at the metric or cost of multipath routes to determine if they are equal cost. All routes to the same destination with different gateways are considered to be equal cost.
- A maximum of four ECMP routes are supported per destination.
- With flow-based handling of ECMP routes, load balancing can be achieved only with a large number of flows.
- When a route is added to or removed from an ECMP route group, the existing flows might take a different path to the destination. For this reason, it is recommended not to change the number of routes to a particular destination once the routes are in use.

Overview of ECMP settings

Following are the parameters, shown with default values, relevant to ECMP routing configuration:


```
[in IP-GLOBAL:ecmp-options]
ecmp-type = none
source-ip-address = no
destination-ip-address = no
ip-protocol = no
```

Parameter	Setting
ecmp-type	Type of ECMP routing to be performed in the system. Following are valid values: none Disable ECMP routing. This is the default value, to conserve processor memory when ECMP routes are not required. load-balanced Use per-packet load balancing, as described in “Packet-based ECMP” on page 4-15. flow-based Use flow-based ECMP routing, as described in “Flow-based ECMP” on page 4-15.
source-ip-address	Use the source IP address of a packet header in defining a packet flow (yes or no). With the default value, the source address is not used.
destination-ip-address	Use the destination IP address of a packet header in defining a packet flow (yes or no). With the default value, the destination address is not used.
ip-protocol	Use the IP protocol field of the packet header in defining a packet flow (yes or no). With the default value, the protocol field is not used.

Sample downstream ECMP configuration

Figure 4-1 shows IPoA connections to four DSL CPE. Traffic destined for the ISP POP will be load balanced on these four links.

Figure 4-1. Four DSL links using ECMP routes



Because packet-based load balancing will be used, the default settings in the IP-GLOBAL `ecmp-options` subprofile apply. To configure the sample setup in Figure 4-1, follow these steps:

- 1 Configure the connection profiles for the IPoA connections.

```
admin> new connection
admin> set station = ipoa_6_3
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 3.3.3.1/8
admin> set ip-options local-address = 3.3.3.2/8
admin> set atm-options vci = 40
admin> set atm-options nailed-group = 253
admin> write -f
admin> new connection
admin> set station = ipoa_6_4
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 4.4.4.1/8
admin> set ip-options local-address = 4.4.4.2/8
admin> set atm-options vci = 71
admin> set atm-options nailed-group = 254
admin> write -f
admin> new connection
admin> set station = ipoa_6_7
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 7.7.7.1/8
admin> set ip-options local-address = 7.7.7.2/8
admin> set atm-options nailed-group = 257
admin> write -f
admin> new connection
```

```
admin> set station = ipoa_6_10
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 10.10.10.1/8
admin> set ip-options local-address = 10.10.10.2/8
admin> set atm-options vci = 40
admin> set atm-options nailed-group = 260
admin> write -f
```

- 2 Configure the static routes to the 9.9.9.1 destination. Each route specifies the same destination and a different gateway address. The CPE addresses are the gateway addresses.

```
admin> new ip-route
admin> set name = ecmp
admin> set dest-address = 9.9.9.1/8
admin> set gateway-address = 3.3.3.1
admin> write -f
admin> new ip-route
admin> set name = ecmp-10
admin> set dest-address = 9.9.9.1/8
admin> set gateway-address = 10.10.10.1
admin> write -f
admin> new ip-route
admin> set name = ecmp-4
admin> set dest-address = 9.9.9.1/8
admin> set gateway-address = 4.4.4.1
admin> write -f
admin> new ip-route
admin> set name = ecmp-7
admin> set dest-address = 9.9.9.1/8
admin> set gateway-address = 7.7.7.1
admin> write -f
```

You can verify that the routes have been established by using the `netstat` command. For example, the following command shows all four ECMP routes activated:

```
admin> netstat -nr
```

Destination	Gateway	IF	Flg	Pref	Met	Use	Age
0.0.0.0/0	110.110.110.1	ie0	SGP	60	1	4	98
3.0.0.0/8	3.3.3.1	wan27	rGT	60	1	1	118
3.3.3.1/32	3.3.3.1	wan27	rT	60	1	9	118
3.3.3.2/32	-	local	rT	0	0	1	118
4.0.0.0/8	4.4.4.1	wan28	rGT	60	1	1	113
4.4.4.1/32	4.4.4.1	wan28	rT	60	1	7	113
4.4.4.2/32	-	local	rT	0	0	1	113
7.0.0.0/8	7.7.7.1	wan29	rGT	60	1	1	104

7.7.7.1/32	7.7.7.1	wan29	rT	60	1	5	104
7.7.7.2/32	-	local	rT	0	0	9	104
8.0.0.0/8	50.50.50.2	ie1-233	SG	60	8	1	98
9.0.0.0/8	10.10.10.1	wan30	SGM	60	8	2	98
9.0.0.0/8	7.7.7.1	wan29	SGM	60	8	2	104
9.0.0.0/8	4.4.4.1	wan28	SGM	60	8	2	113
9.0.0.0/8	3.3.3.1	wan27	SGM	60	8	1	118
10.0.0.0/8	10.10.10.1	wan30	rGT	60	1	1	98
10.10.10.1/32	10.10.10.1	wan30	rT	60	1	3	98
10.10.10.2/32	-	local	rT	0	0	1	98
50.0.0.0/8	-	ie1-233	C	0	0	15	51696
50.50.50.1/32	-	local	CP	0	0	1	51696
100.0.0.0/8	-	ie1	C	0	0	1	51696
100.1.1.1/32	-	local	CP	0	0	1	51696
110.0.0.0/8	-	ie0	C	0	0	64597	57440
110.110.110.71/32	-	local	CP	0	0	23358	57440
127.0.0.0/8	-	bh0	CP	0	0	0	57440
127.0.0.1/32	-	local	CP	0	0	0	57440
127.0.0.2/32	-	rj0	CP	0	0	0	57440
224.0.0.0/4	-	mcast	CP	0	0	0	57440
224.0.0.1/32	-	local	CP	0	0	109	57440
224.0.0.2/32	-	local	CP	0	0	0	57440
224.0.0.9/32	-	local	CP	0	0	0	57440
224.0.0.13/32	-	local	CP	0	0	0	57440
255.255.255.255/32	-	ie0	CP	0	0	0	57440
Total Routes = 33 Hidden Routes = 0							

You can verify load distribution on the DSL ports by using the atmvcmmgr command. For example, the following command on the SDSL module shows load balancing achieved with all four routes up.

```
sdsl-atm-v2-1/6> atmvcmmgr -t 6
```

port/vpi/vci	TX Cells	Discarded Cells	TX packets	Discarded packets	Error packets	Cells in last Pkt
3/ 0/ 40	9	0	0	0	0	3
4/ 0/ 71	12	0	0	0	0	3
7/ 0/ 35	12	0	0	0	0	3
10/ 0/ 40	12	0	0	0	0	3

```
sdsl-atm-v2-1/6> atmvcmmgr -t 6
```

port/vpi/vci	TX Cells	Discarded Cells	TX packets	Discarded packets	Error packets	Cells in last Pkt
3/ 0/ 40	12	0	0	0	0	3
4/ 0/ 71	12	0	0	0	0	3
7/ 0/ 35	12	0	0	0	0	3
10/ 0/ 40	12	0	0	0	0	3

Example of downstream flow-based ECMP

In this example, the same physical network and the same CONNECTION and IP-ROUTE profiles used in “Sample downstream ECMP configuration” on page 4-17. This example shows how to modify the setup for flow-based ECMP routing on the basis of the destination address in traffic streams forwarded to the 9.0.0.0 network through the four DSL links.

For example, suppose 15 traffic streams are being forward in the ECMP routes, to 15 different destinations in the 9.0.0.0 network in Figure 4-1 (page 4-17). The packets belonging to each stream will always take the same path across the same DSL link.

After configuring the connection and IP-ROUTE profiles as described in “Sample downstream ECMP configuration” on page 4-17, configure flow-based ECMP routing as follows:

```
admin> read ip-global
admin> set ecmp-options ecmp-type = flow-based
admin> set ecmp-options destination-ip-address = yes
admin> write -f
```

The following command shows all four ECMP routes activated:

```
admin> netstat -nr
```

Destination	Gateway	IF	Flg	Pref	Met	Use	Age
0.0.0.0/0	110.110.110.1	ie0	SGP	60	1	2	428
3.0.0.0/8	3.3.3.1	wan23	rGT	60	1	1	690
3.3.3.1/32	3.3.3.1	wan23	rT	60	1	81	690
3.3.3.2/32	-	local	rT	0	0	43	690
4.0.0.0/8	4.4.4.1	wan24	rGT	60	1	1	606
4.4.4.1/32	4.4.4.1	wan24	rT	60	1	26	606
4.4.4.2/32	-	local	rT	0	0	9	606
7.0.0.0/8	7.7.7.1	wan25	rGT	60	1	1	554
7.7.7.1/32	7.7.7.1	wan25	rT	60	1	13	554
7.7.7.2/32	-	local	rT	0	0	39	554
8.0.0.0/8	50.50.50.2	ie1-233	SG	60	8	1	428
9.0.0.0/8	10.10.10.1	wan26	SGM	60	8	2	499
9.0.0.0/8	7.7.7.1	wan25	SGM	60	8	4	554
9.0.0.0/8	4.4.4.1	wan24	SGM	60	8	4	606
9.0.0.0/8	3.3.3.1	wan23	SGM	60	8	4	690
10.0.0.0/8	10.10.10.1	wan26	rGT	60	1	1	499
10.10.10.1/32	10.10.10.1	wan26	rT	60	1	10	499
10.10.10.2/32	-	local	rT	0	0	3	499
50.0.0.0/8	-	ie1-233	C	0	0	3	428
50.50.50.1/32	-	local	CP	0	0	1	428
100.0.0.0/8	-	ie1	C	0	0	1	428
100.1.1.1/32	-	local	CP	0	0	1	428
110.0.0.0/8	-	ie0	C	0	0	47061	6172
110.110.110.71/32	-	local	CP	0	0	17538	6172
127.0.0.0/8	-	bh0	CP	0	0	0	6172
127.0.0.1/32	-	local	CP	0	0	0	6172
127.0.0.2/32	-	rj0	CP	0	0	0	6172
224.0.0.0/4	-	mcast	CP	0	0	0	6172
224.0.0.1/32	-	local	CP	0	0	12	6172
224.0.0.2/32	-	local	CP	0	0	0	6172
224.0.0.9/32	-	local	CP	0	0	0	6172
224.0.0.13/32	-	local	CP	0	0	0	6172
255.255.255.255/32	-	ie0	CP	0	0	0	6172
Total Routes = 33		Hidden Routes = 0					

The traffic distribution will indicate flow-based routing, For example:


```
sdsl-atm-v2-1/6> atmccmgr -t 6
```

port/vpi/vci	TX Cells	Discarded Cells	TX packets	Discarded packets	Error packets	Cells in last Pkt
3/ 0/ 40	681	0	0	0	0	3
4/ 0/ 71	846	0	0	0	0	3
7/ 0/ 35	672	0	0	0	0	3
10/ 0/ 40	675	0	0	0	0	3

Sample upstream ECMP through routed VLAN interfaces

Figure 4-2 shows two upstream ECMP routes to the same destination system. The routes use routed VLAN interfaces on a Gigabit Ethernet port.

Figure 4-2. Two routed VLAN uplinks using ECMP routes



In this example, packet-based load balancing will be used, so the default settings in the IP-GLOBAL `ecmp-options` subprofile apply. To configure the sample setup in Figure 4-1, follow these steps:

- 1 Verify that the default settings are in effect in `ecmp-options`. For example:

```
admin> get ip-global ecmp-options
[in IP-GLOBAL:ecmp-options]
ecmp-type = load-balanced
source-ip-address = no
destination-ip-address = no
admin> ip-protocol = no
```

- 2 Configure the CONNECTION profile to the downstream CPE. For example:

```
admin> new connection
admin> set station = ipoa_6_3
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 3.3.3.1/8
admin> set ip-options local-address = 3.3.3.2/8
admin> set atm-options vci = 40
admin> set atm-options nailed-group = 253
admin> write -f
```

- 3 Configure the two routed VLAN interfaces.

```
admin> new vlan { { shelf-1 second-control-module 2 } 678 }
admin> set enabled = yes
admin> write -f
admin> new ip-interface { { shelf-1 second-control-module 2 } 678 }
```

```

admin> set ip-address = 60.60.60.1/8
admin> set vlan-enabled = yes
admin> set vlan-id = 678
admin> write -f
admin> new vlan { { shelf-1 second-control-module 2 } 233 }
admin> set enabled = yes
admin> write -f
admin> new ip-interface { { shelf-1 second-control-module 2 } 233 }
admin> set ip-address = 50.50.50.1/8
admin> set vlan-enabled = yes
admin> set vlan-id = 233
admin> write -f

```

4 Configure the ECMP routes.

```

admin> new ip-route g1
admin> set dest-address = 8.1.1.1/8
admin> set gateway-address = 50.50.50.2
admin> write -f
admin> new ip-route g2
admin> set dest-address = 8.1.1.1/8
admin> set gateway-address = 60.60.60.2
admin> write -f

```

Following is the output of the netstat -nr command relevant to the routed VLAN ECMP configuration:

```

.....
8.0.0.0/8          60.60.60.2      ie1-678      SGM      60    8      29
3498
8.0.0.0/8          50.50.50.2      ie1-233      SGM      60    8      13
3498
.....
50.0.0.0/8        -                ie1-233      C         0    0      13      3683
50.50.50.1/32     -                local        CP        0    0      13      3683
60.0.0.0/8        -                ie1-678      C         0    0      15      3683
60.60.60.1/32     -                local        CP        0    0      15      3683
.....

```

NP's routing table:

Route Table for vrouter:main

Destination	Gateway	IF	DID	needArp
3.0.0.0/8	3.3.3.1	wan22	045012	N
3.3.3.1/32	3.3.3.1	wan22	045012	N
8.0.0.0/8	50.50.50.2	ie1-233	045013	N
8.0.0.0/8	60.60.60.2	ie1-678	045014	N
50.50.50.2/32	50.50.50.2	ie1-233	045013	N
60.60.60.2/32	60.60.60.2	ie1-678	045014	N

Total 39 routes <end>

Configuring and using address pools

An address pool is a range of contiguous addresses on a local IP network or subnet. Pool addresses are available for assignment to incoming connections that request an address. When the call terminates, the address is returned to the pool, making it available again for assignment.

Typically, a user's CONNECTION profile specifies an address-pool index to request an address within a particular range, or specifies the zero index to indicate that it will accept the next available address from any pool.

For increased control over the user address space, you can map an address-pool index to a specific LIM slot, or to a specific LIM slot and virtual router combination, as described in "Slot-based address assignment" on page 4-36.

Pool requirements

If you designate a subnet for IP address pools, you must make sure that other IP hosts on the local network know the route to that subnet. You must also make sure that the pools do not overlap (that the pools do not contain duplicate addresses).

For related information, see "Defining address pools for a virtual router" on page 7-7.

Overview of settings for defining pools

You can define up to 128 address pools locally in the IP-GLOBAL profile. Those pools can be used to assign addresses to connections authenticated locally (in CONNECTION profiles) or by RADIUS. If you are using RADIUS authentication, you can choose to define address pools in RADIUS instead of, or in addition to, those defined locally. If you have the RADIPAD program installed, you can use it to manage address pools centrally on a single RADIUS server.

Settings in the IP-GLOBAL profile

The following parameters (shown with default values) configure address pools locally:

```
[in IP-GLOBAL]
pool-summary = no
pool-base-address = [ 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0+
assign-count = [ 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 +
pool-name = [ "" "" "" "" "" "" "" "" "" "" "" "" "" "" "" "" "" "" "" "" +
must-accept-address-assign = no
```

Parameter	Setting
pool-summary	Set/clear the Pool Summary flag. For details, see "Example of configuring summarized address pools" on page 4-27.
pool-base-address	Base address of a pool of contiguous addresses on a local network or subnet.
assign-count	Number of addresses in the pool.

Parameter	Setting
pool-name	A pool name, required only when TACACS+ authentication is in use. If TACACS+ authentication is not in use, the name is treated as a comment.
must-accept-address-assign	Enables or disables rejection of an assigned IP address by an incoming caller during PPP negotiation.

Settings in RADIUS pseudo-user profiles

You can define address pools in a RADIUS pools pseudo-user profile. The first line of pools pseudo-user profile uses the following format:

```
pools-name Password = "ascend", Service-Type = Outbound-User
```

The `name` argument is the system name (specified by the `name` parameter in the SYSTEM profile). Subsequent lines in the profile define IP address pools by using the Ascend-IP-Pool-Definition (217) attribute. The value of the Ascend-IP-Pool-Definition attribute uses the following syntax:

```
"pool-num base-addr assign-count"
```

Syntax element	Description
pool-num	Pool number. If you use the same number to designate two pools, one locally and one in RADIUS, the RADIUS definition takes precedence. So if you have defined some pools in the IP-GLOBAL profile and do not wish to override them, start numbering the pools at the next number. For example, if you defined 10 pools in the IP-GLOBAL profile, start with number 11 in RADIUS. Otherwise, start with 1.
base-addr	The base address in a pool of contiguous addresses on the local network or subnet.
assign-count	Number of addresses included in the pool.

Global RADIUS pools (RADIPAD)

RADIUS IP Address Daemon (RADIPAD) is a program that works with RADIUS authentication to manage IP address pools centrally, so that connections can all acquire an address from a global pool, regardless of which system answers the call.

RADIPAD runs on one RADIUS server on the network. A Stinger unit sends an authentication request to RADIUS, and if the user profile contains an attribute to allocate an IP address from the global pool, RADIUS sends a request to RADIPAD to acquire the address.

The Stinger unit does not communicate directly with RADIPAD, so it does not require additional configuration to use RADIPAD. To configure RADIPAD, you define the global pools of addresses, specify which RADIUS server is running RADIPAD, and (optionally) specify which systems can obtain addresses from those pools. You can then create RADIUS user profiles that acquire an IP address from the global pool.

At startup, syslog notes RADIUS requests to release RADIUS-allocated IP addresses. Some versions of the RADIUS server might time out the request, resulting in log messages indicating the release of global-pool addresses.

Defining global pools

Global address pools are defined in a `global-pools` pseudo-user profile on the server running RADIPAD. The first line of a `global-pools` pseudo-user profile uses the following format:

```
global-pools-name Password = "ascend", Service-Type = Outbound-User
```

The *name* argument is a designation for any class of users. You can create multiple global pool profiles for multiple user classes. For example, you could create profiles named `global-pool-ppp`, `global-pool-slip`, and so forth. Subsequent lines in the profile define IP address pools by using the `Ascend-IP-Pool-Definition (217)` attribute. This is the same attribute described in “Settings in RADIUS pseudo-user profiles” on page 4-24, and it follows the same rules for global pools. In addition, when the Stinger unit assigns an address from a pool managed by the RADIPAD daemon, RADIPAD tries to allocate an address from the pools in order, by pool number, and chooses an address from the first pool with an available IP address.

Specifying the RADIPAD host

Each RADIUS server must specify the host running RADIPAD and (optionally) the systems that can access the global pools. These settings are defined in a `radipa-hosts` pseudo-user profile, which uses the following format in the first line of the profile:

```
radipa-hosts Password = "ascend", Service-Type = Outbound-User
```

Subsequent lines in the profile use the following attribute-value pairs to define which hosts can assign addresses from the pools managed by RADIPAD:

RADIUS attribute	Value
Ascend-Assign-IP-Client (144)	Address of a system that is allowed to access the global address pools managed by RADIPAD. You can specify multiple instances of this attribute. If no client addresses are specified, all units listed in the RADIUS clients file can access RADIPAD pools.
Ascend-Assign-IP-Server (145)	Address of the server running RADIPAD. Only one instance of this attribute can appear in the profile, and it must specify the correct IP address.

For example:

```
radipa-hosts Password = "ascend", Service-Type = Outbound-User
  Ascend-Assign-IP-Server = 10.31.4.34,
  Ascend-Assign-IP-Client = 10.31.4.10,
  Ascend-Assign-IP-Client = 10.31.4.11
```

You can specify only one RADIPAD server, but you can include multiple clients. The sample profile indicates that two systems (10.31.4.10 and 10.31.4.11) can access RADIPAD pools as clients.

Preventing the use of class boundary addresses

If you define address pools that contain more than 254 addresses, be aware that the system allocates the class boundary addresses (*n.n.n.0* and *n.n.n.255*) as valid connection addresses. According to CIDR, this is permitted because the pool is not a /24 network. However, some client systems, such as Windows, do not tolerate the class boundary addresses well. For example, because Windows assumes a /24 network, it broadcasts NetBIOS over IP name service to the .255 address, which could overwhelm a connection assigned the .255 host address.

To prevent client software from using a host address for broadcasts, you must explicitly apply a filter that prevents the system from using the class boundary addresses. For example, if you are using RADIUS authentication, you can apply a data filter, in the ANSWER-DEFAULTS profile, that drops packets from any source to pool address *n.n.n.0* or *n.n.n.255*.

Examples of configuring address pools

For a pool that is not summarized, each assigned address is advertised as its own host route. Such a pool can start at any base address. Addresses do not accept a subnet mask component, because they are always advertised as host routes.

The following commands define three address pools, each containing 50 addresses. Pool 1 contains 10.2.3.4 through 10.2.3.54. Pool 2 contains 11.5.7.51 through 11.5.7.101. Pool 3 contains 12.7.112.15 through 12.7.112.65.

```
admin> read ip-global
admin> set pool-base-address 1 = 10.2.3.4
admin> set pool-base-address 2 = 11.5.7.51
admin> set pool-base-address 3 = 12.7.112.15
admin> set assign-count 1 = 50
admin> set assign-count 2 = 50
admin> set assign-count 3 = 50
admin> write -f
```

Following is a comparable RADIUS pools profile (for use by a single RADIUS server):

```
pools-taos01 Password = "ascend", Service-Type = Outbound-User
  Ascend-IP-Pool-Definition = "1 10.2.3.4 50",
  Ascend-IP-Pool-Definition = "2 11.5.7.51 50",
  Ascend-IP-Pool-Definition = "3 12.7.112.15 50"
```

Following is a comparable global pools definition (for use with RADIPAD):

```
global-pool-ppp Password = "ascend", Service-Type = Outbound-User
  Ascend-IP-Pool-Definition = "1 10.2.3.4 50",
  Ascend-IP-Pool-Definition = "2 11.5.7.51 50",
  Ascend-IP-Pool-Definition = "3 12.7.112.15 50"
```

Although some client software assumes a default subnet mask of 255.255.255.0 for PPP interfaces, you can define pools on subnets wider than /24. For example, the following commands define an address pool on a /23 subnet:

```
admin> read ip-global
admin> set pool-base-address 1 = 10.55.178.1
```

```
admin> set assign-count 1 = 510
admin> write -f
```

This pool definition translates to 10.55.178.0/23 (a subnet mask of 255.255.254.0). Following are comparable RADIUS definitions:

```
pools-taos01 Password = "ascend", Service-Type = Outbound-User
Ascend-IP-Pool-Definition = "1 10.55.178.1 510"

global-pool-ppp Password = "ascend", Service-Type = Outbound-User
Ascend-IP-Pool-Definition = "1 10.55.178.1 510"
```

Example of configuring summarized address pools

The `pool-summary` feature reduces routing overhead associated with address pools. To use this feature, you must do the following:

- Set `suppress-host-routes` to `yes`, as described in “Suppressing host-route advertisements” on page 4-13.
- Set `pool-summary` to `yes` in the IP-GLOBAL profile.
- Define pools that are network aligned, as described in “Defining network-aligned pools” on page 4-28.

Overview of summary routes

A *summary route* is a network-level route that the system advertises instead of advertising each individual host route as it becomes active. The system suppresses the host route advertisements and instead advertises a summary route to the network-aligned pool. The summary route is used only for advertising purposes, to reduce the overhead of routing protocols.

For example, with the following configuration:

```
admin> read ip-global
admin> set pool-summary = yes
admin> set pool-base-address 1 = 195.195.195.1
admin> set assign-count 1 = 254
admin> write -f
```

The system creates the following summary route, which points to the “reject” (rj0) interface:

Destination	Gateway	IF	Flg	Pref	Met	Use	Age
195.195.195.0/24	-	rj0	C	0	0	0	523510

When connections come up using this pool, the system adds a host route for each connection, but it flags the host route as “private” (the P flag). For example:

Destination	Gateway	IF	Flg	Pref	Met	Use	Age
195.195.195.0/24	-	rj0	C	0	0	0	523875
195.195.195.1/32	195.195.195.1/32	wan15	rPT	60	1	1	323
195.195.195.5/32	195.195.195.5/32	wan23	rPT	60	1	1	127

With these routes, the Stinger system advertises only the summary route. Another router receiving the advertisement interprets the summary route as an ordinary network route, and forwards traffic destined for 195.195.195.0/24 to the Stinger system. When the system receives the traffic, it forwards traffic addressed to

195.195.195.1 and 195.195.195.5 to the appropriate WAN interfaces. All other traffic to 195.195.195.0/24 is directed to the rj0 interface (it is discarded).

Setting the pool-summary flag

The following commands enable the pool-summary flag:

```
admin> read ip-global
admin> set pool-summary = yes
admin> write -f
```

Defining network-aligned pools

Following are the rules for network-aligned address pools:

- The specified number of addresses in the pool must be two less than the total number of addresses in the pool. (Add 2 to the assign-count value for the total number of addresses in the subnet, and calculate the mask for the subnet on the basis of this total.)
$$\text{assign-count} + 2 = \text{number of subnet hosts}$$
- The specified base address of the pool must be the first host address. (Subtract 1 from the pool-base-address value for the base address for the subnet.)
$$\text{pool-base-address} - 1 = \text{network-aligned subnet address}$$

The following commands enable the pool-summary flag and define a network-aligned pool:

```
admin> read ip-global
admin> set pool-summary = yes
admin> set assign-count 1 = 62
admin> set pool-base-address 1 = 10.12.253.1
admin> write -f
```

In the preceding sample configurations, the assign-count parameter is set to 62. When you add 2 to this value, you get 64. The subnet mask for 64 addresses is 255.255.255.192 ($256 - 64 = 192$). The prefix length for a 255.255.255.192 mask is /26.

The pool-base-address parameter is set to 10.12.253.1. When you subtract 1 from this value, you get 10.12.253.0, which is a valid network-aligned base address for the 255.255.255.192 subnet mask. (Note that 10.12.253.64, 10.12.253.128, and 10.12.253.192 are also valid zero addresses for the same mask.) The resulting address pool subnet is 10.12.253.0/26.

Following is a comparable RADIUS pools profile (for use by a single RADIUS server).

```
pools-taos01 Password = "ascend", Service-Type = Outbound-User
Ascend-IP-Pool-Definition = "1 10.12.253.1 62"
```

Following is a comparable global pools definition (for use with RADIPAD):

```
global-pool-ppp Password = "ascend", Service-Type = Outbound-User
Ascend-IP-Pool-Definition = "1 10.12.253.1 62"
```

The system still creates (but does not advertise) a host route for each assigned address in the pool. Host routes take precedence over subnet routes, so packets whose

destination matches an assigned IP address from the pool are routed properly. However, because the system advertises the entire pool as a route, and only privately knows which IP addresses in the pool are active, a remote network might improperly send the Stinger unit a packet for an inactive IP address. If that occurs, the packets are routed to the Reject (rj0) interface (127.0.0.2). Packets routed to the Reject interface are bounced back to the sender with an ICMP unreachable message.


Examples of assigning an address from a pool

When an incoming call requests an IP address, the Stinger unit assigns one from a pool. A host requests an address if its client software has settings such as the following:

```
Username=victor
Accept Assigned IP=Yes
IP address=Dynamic (or Assigned or N/A)
Netmask=255.255.255.255 (or None or N/A)
Default Gateway=None or N/A
Name Server=10.2.3.55
Domain suffix=abc.com
Baud rate=38400
Hardware handshaking ON
VAN Jacobson compression ON
```

Figure 4-3 shows a remote host requesting and being assigned an IP address.

Figure 4-3. Remote CPE requiring dynamic IP address assignment



The following commands enable dynamic address assignment systemwide:

```
admin> read answer-defaults
admin> set ip-answer assign-address = yes
admin> write -f
```

During PPP negotiation, a CPE can reject an IP address offered by the router and present the caller's own IP address for consideration. For security purposes, many sites set `must-accept-address-assign` to `yes` to ensure that the Stinger unit terminates such a call, as shown in the following example:

```
admin> read ip-global
admin> set must-accept-address-assign = yes
admin> write -f
```

For address assignment to occur, the Stinger unit must have an address available for assignment, the ANSWER-DEFAULTS profile must enable dynamic assignment, the

client profile must specify dynamic assignment, and the client's PPP software must be configured to acquire its IP address dynamically.

The following commands configure a profile to acquire an address from the first pool that has available addresses:

```
admin> new connection victor
admin> set active = yes
admin> set encapsulation-protocol = ppp
admin> set ppp rcv-password = localpw
admin> set ip-options address-pool = 0
admin> write -f
```

Following is a comparable RADIUS profile:

```
victor Password = "localpw"
    Service-Type = Framed-User,
    Framed-Protocol = PPP,
    Ascend-Assign-IP-Pool = 0
```

Following is a comparable RADIUS profile that acquires an address from any global pool managed by the RADIPAD daemon:

```
victor Password = "localpw"
    Service-Type = Framed-User,
    Framed-Protocol = PPP,
    Ascend-Assign-IP-Pool = 65535,
    Ascend-Assign-ip-global-Pool = "global-pool-ppp"
```

IP pool chaining


Because the addresses within a pool must be contiguous, many sites have defined a large number of pools, with each pool containing only a small range of addresses. For example, the following RADIUS profile defines six pools, each containing 10 addresses:

```
pools-JFAN-TNT Password = "ascend"
    Service-Type = Outbound,
    Ascend-IP-Pool-Definition = "1 11.168.6.10 10",
    Ascend-IP-Pool-Definition = "2 12.168.6.10 10",
    Ascend-IP-Pool-Definition = "3 13.168.6.10 10",
    Ascend-IP-Pool-Definition = "7 17.168.6.10 10",
    Ascend-IP-Pool-Definition = "8 18.168.6.10 10",
    Ascend-IP-Pool-Definition = "9 19.168.6.10 10"
```

In earlier versions of the software, you could allow a connection to acquire an address from any pool (by assigning the pool number 0 in the connection's profile) or from a single specified pool, such as pool 1. IP pool chaining enables you to allow a connection to acquire an address from any pool within a chain.

When IP pool chaining is enabled, contiguous pools are treated as one *pool space* with shared addresses. When the system assigns an address to an end user, it begins searching for an available address in the first pool of the chain and stops when it either finds an available address or encounters a null pool definition. So, the pools within a chain must be defined in a contiguous sequence. For example, the following


```
pools-JFAN-TNT Password = "ascend", Service-Type = Outbound
Ascend-IP-Pool-Chaining = IP-Pool-Chaining-Yes,
Ascend-IP-Pool-Definition = "1 11.168.6.10 10",
Ascend-IP-Pool-Definition = "2 12.168.6.10 10",
Ascend-IP-Pool-Definition = "3 13.168.6.10 10",
Ascend-IP-Pool-Definition = "7 17.168.6.10 10",
Ascend-IP-Pool-Definition = "8 18.168.6.10 10",
Ascend-IP-Pool-Definition = "9 19.168.6.10 10"
```



```
[in IP-GLOBAL]
pool-chaining = yes
pool-base-address = [ 172.20.31.1 172.20.33.1 0.0.0.0 153.37.21.1 0.0+
assign-count = [ 126 126 0 30 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 +
```

Whether pool chains are defined locally or in RADIUS, the pool addresses are available for dynamic assignment regardless of where the connection's profile is authenticated.

Following are the parameters, shown with default settings, relevant to IP pool chaining:

```
[in IP-GLOBAL]
pool-chaining = no
pool-base-address = [ 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0+
assign-count = [ 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0+
[in CONNECTION/"" :ip-options]
address-pool = 0
```

4-31

Parameter	Setting
pool-base-address	An array of up to 128 IP addresses to be used as the first address in a pool. These values are used with the assign-count values to define address pools locally. A pool chain contains all of the pools defined in sequence within the array, such as 1, 2, 3. To end a pool chain, leave a null value in the array.
assign-count	An array of up to 128 numbers that specify the number of addresses in a pool that starts with the corresponding pool-base-address.
address-pool	Number of an address pool from which to acquire an address. When pool chaining is enabled, a pool number within a chain includes addresses defined in all other pools within the chain. For example, if pools 1, 2, and 3 are in a pool chain, setting this parameter to 1 has the same effect as setting it to 2 or 3.

Example of local pool chain definition

The following commands define five address pools, which form two pool chains. Notice that the pool numbers (their indexes in the pool-base-address and assign-count arrays) are contiguous within a chain.

```
admin> read ip-global
admin> set pool-chaining = yes
admin> set pool-base-address 1 = 10.1.1.1
admin> set pool-base-address 2 = 11.1.1.1
admin> set pool-base-address 3 = 12.1.1.1
admin> set assign-count 1 = 50
admin> set assign-count 2 = 50
admin> set assign-count 3 = 50
admin> set pool-base-address 7 = 13.1.1.1
admin> set pool-base-address 8 = 14.1.1.1
admin> set assign-count 7 = 50
admin> set assign-count 8 = 50
admin> write -f
```

The following commands enable dynamic address assignment systemwide:

```
admin> read answer-defaults
admin> set ip-answer assign = yes
admin> write -f
```

The following commands configure profiles to acquire an address from the first pool chain. When the end users initiate a session request, they can acquire an address from 10.1.1.1 to 10.1.1.51, from 11.1.1.1 to 11.1.1.51, or from 12.1.1.1 to 12.1.1.51. If no addresses are available within those ranges, the connection is refused.

```
admin> new connection jfan
admin> set active = yes
admin> set encapsulation-protocol = ppp
admin> set ppp-options recv-password = localpw
admin> set ip-options address-pool = 2
admin> write -f
admin> new connection ravi
admin> set active = yes
admin> set encapsulation-protocol = ppp
admin> set ppp-options recv-password = localpw
admin> set ip-options address-pool = 1
admin> write -f
```

Following are comparable RADIUS profiles:

```
jfan Password = "localpw"
    Service-Type = Framed-User,
    Framed-Protocol = PPP,
    Ascend-Assign-IP-Pool = 2
ravi Password = "localpw"
    Service-Type = Framed-User,
    Framed-Protocol = PPP,
    Ascend-Assign-IP-Pool = 1
```

Pool chaining in RADIUS

Whether pool chains are defined locally or in a RADIUS pool's pseudo-user profile, the pool addresses are available for dynamic assignment regardless of where the connection's profile is authenticated.

Overview of RADIUS profile settings

RADIUS servers use the following attribute-value pairs to define and apply pool chains:

RADIUS attribute	Value
Ascend-IP-Pool-Chaining (85)	Enables or disables IP pool chaining in a pseudo-user profile that defines address pools. If this attribute is set to IP-Pool-Chaining-Yes (1), the system treats contiguous IP address pools as a single extended pool space when searching for an available address to assign to a connection. With a value of IP-Pool-Chaining-No (0), the system treats each address pool as a separate space. Note When this attribute is specified in a RADIUS profile, its value overrides the Pool-Chaining setting in the IP-GLOBAL profile.
Ascend-IP-Pool-Definition (217)	Address pool definition in a pseudo-user profile. The value has the following syntax: <code>pool-number base-addr assign-count</code> The <code>pool-number</code> value is an integer that identifies the pool. A pool chain contains all of the pools defined in sequence, such as 1, 2, 3. To end a pool chain, leave a gap in the sequence of <code>pool-number</code> values. The <code>base-addr</code> value is an IP address to be used as the first address in a pool, and the <code>assign-count</code> value specifies the number of addresses in a pool.
Ascend-Assign-IP-Pool (218)	Number of the address pool from which the RADIUS user profile should acquire an address. When pool chaining is enabled, a pool number within a chain includes addresses defined in all other pools within the chain. For example, if pools 1, 2, and 3 are in a pool chain, setting this value to 1 has the same effect as setting it to 2 or 3.

To use these attributes, the RADIUS server must support vendor-specific attributes (VSAs) and the system must be configured in VSA compatibility mode. Following are the relevant settings:

```
[in EXTERNAL-AUTH]
auth-type = radius

[in EXTERNAL-AUTH:rad-auth-client]
auth-radius-compatible = vendor-specific
```

For details about these settings, see the *Stinger Reference*.

Example of pool chaining in RADIUS

The following pseudo-user profile defines five address pools, which form two pool chains. Notice that the pool numbers are contiguous within a chain.

```
pools-JFAN-TNT Password = "ascend"
Service-Type = Outbound,
  Ascend-IP-Pool-Chaining = IP-Pool-Chaining-Yes,
  Ascend-IP-Pool-Definition = "1 10.1.1.1 50",
  Ascend-IP-Pool-Definition = "2 11.1.1.1 50",
  Ascend-IP-Pool-Definition = "3 12.1.1.1 50",
  Ascend-IP-Pool-Definition = "7 13.1.1.1 50",
  Ascend-IP-Pool-Definition = "8 14.1.1.1 50"
```

The following commands configure local CONNECTION profiles to acquire an address from the first pool chain. When the end users initiate a session request, they can acquire an address from 13.1.1.1 to 13.1.1.51, or from 14.1.1.1 to 14.1.1.51. If no addresses are available within those ranges, the connection is refused.

```
admin> new connection hanif
admin> set active = yes
admin> set encapsulation-protocol = ppp
admin> set ppp-options recv-password = localpw
admin> set ip-options address-pool = 7
admin> write -f
admin> new connection hasnain
admin> set active = yes
admin> set encapsulation-protocol = ppp
admin> set ppp-options recv-password = localpw
admin> set ip-options address-pool = 8
admin> write -f
```

Following are comparable RADIUS user profiles:

```
hanif Password = "localpw"
  Service-Type = Framed-User,
  Framed-Protocol = PPP,
  Ascend-Assign-IP-Pool = 7
hasnain Password = "localpw"
  Service-Type = Framed-User,
  Framed-Protocol = PPP,
  Ascend-Assign-IP-Pool = 8
```

Slot-based address assignment

When you use slot-based address assignment, all connections established on a slot obtain an address from the same pool. If the slot contains a Stinger GE-OLIM, all connections on all remote LIMs connected to that GE-OLIM obtain an address from the same pool.



Note Slot-based address assignment causes the system to select an address-pool index for an incoming connection from a configured `slot-vr-config` profile rather than from the user's `connection` profile. Once the system has selected the address-pool index, however, the rest of the system behavior for address pool capabilities, including pool chaining, the `share-global-pool` feature for virtual routers, and so forth, still operate in the same way.

For address assignment to occur, the Stinger system must have an address available for assignment, the `answer-defaults` profile must enable dynamic assignment, the client profile must specify dynamic assignment, and the client's PPP software must be configured to acquire its IP address dynamically.

Overview of profile settings

Following are the profile settings, shown with default values, for slot-based address assignment:

```
[in IP-GLOBAL]
pool-selection = connection-based

[in VROUTER/"" ]
pool-selection = connection-based

[in SLOT-VR-CONFIG/{ { any-shelf any-slot 0 } "" }]
index* = { { any-shelf any-slot 0 } "" }
address-pool = 0
```


Parameter	Setting
pool-selection	Specifies the method by which to determine the address-pool index to use in selecting a pool.
	connection-based Select a pool on the basis of the address-pool setting of the subscriber's connection profile. This is the default setting. slot-based Select a pool on the basis of the address-pool setting of the slot-vr-config profile for the slot on which the connection is being established. If the slot-vr-config profile is not present, address allocation fails and the connection is not established.

Parameter	Setting				
index	Specifies the physical address of a LIM or GE-OLIM slot and optionally, a virtual router name. Slot-based address assignment is not enabled for IMA LIMs. Any virtual router for which the <code>pool-selection</code> parameter is set to <code>slot-based</code> (including the global router) uses the <code>slot-vr-config</code> index to select an address pool for incoming connections. The index consists of the following fields: <table><tr><td>physical-address</td><td>The <code>physical-address</code> field specifies the shelf number (which must be 1), the LIM slot number, and a 0 item number (which indicates the whole slot).</td></tr><tr><td>vrrouter</td><td> The default null string represents the global router. This is the only supported value for this software version. If you enter any other string of characters, it must represent the index of a configured VROUTER profile. Do not enter the <code>global-vrouter</code> string from the <code>ip-global</code> profile to represent the global router. </td></tr></table>	physical-address	The <code>physical-address</code> field specifies the shelf number (which must be 1), the LIM slot number, and a 0 item number (which indicates the whole slot).	vrrouter	The default null string represents the global router. This is the only supported value for this software version. If you enter any other string of characters, it must represent the index of a configured VROUTER profile. Do not enter the <code>global-vrouter</code> string from the <code>ip-global</code> profile to represent the global router.
physical-address	The <code>physical-address</code> field specifies the shelf number (which must be 1), the LIM slot number, and a 0 item number (which indicates the whole slot).				
vrrouter	The default null string represents the global router. This is the only supported value for this software version. If you enter any other string of characters, it must represent the index of a configured VROUTER profile. Do not enter the <code>global-vrouter</code> string from the <code>ip-global</code> profile to represent the global router.				
address-pool	Specifies the index of an address pool (a number from 0 to 150) to be used for connections established on the slot/virtual router combination identified in the <code>index</code> parameter. Multiple slots can share an address pool by configuring multiple <code>slot-vr-config</code> profiles with different <code>index</code> values and the same <code>address-pool</code> .				

Sample slot-based address configuration using the global router

Figure 4-4 shows a Stinger IP DSLAM with a CPE configured as a bridge to send PPPoE packets. The WAN interface of the CPE connects to a LIM in slot 2 of the Stinger.

Figure 4-4. Example of a PPPoE session using slot-based address assignment



The CPE has a configured profile used to bridge IP packets between Ethernet and WAN interfaces, in which dynamic address assignment is enabled. To configure address pools and a CPE connection that uses slot-based dynamic address assignment, follow these steps:

- 1 Configure address pools on the host. For example:
admin> read ip-global
admin> set pool-summary = yes
admin> set pool-selection = slot-based
admin> set pool-base-address 11 = 194.194.1.1
admin> set pool-base-address 12 = 194.194.2.1
admin> set pool-base-address 13 = 194.194.3.1
admin> set assign-count 11 = 254
admin> set assign-count 12 = 254
admin> set assign-count 13 = 254
admin> write -f
- 2 Configure the system-generated slot-vr-config profile for LIM slot 2. Specify one of the configured address pools for use by this slot. For example:
admin> read slot-vr-config { { 1 2 0 } }
admin> set address-pool = 12
admin> write -f
- 3 Enable PPP authentication and dynamic address assignment in the answer-defaults profile.
admin> read answer-defaults
admin> set ppp-answer receive-auth-mode = any-ppp-auth
admin> set ppp-answer bi-directional-auth = allowed
admin> set ip-answer assign-address = yes
admin> write -f
- 4 Create the ATM circuit connection profile that provides the interface over which ATM frames are received from the CPE. This profile terminates on the LIM internal port.
admin> new connection
admin> set station = atmckt26
admin> set active = yes
admin> set mp-options enabled = no
admin> set atm-options vci = 37
admin> set atm-options nailed-group = 51
admin> set atm-connect-options vci = 37
admin> set atm-connect-options nailed-group = 99
admin> write -f
- 5 Configure the PPPoE connection profile used to establish PPP connectivity. For example:
admin> new connection
admin> set station = pppoe2
admin> set active = yes
admin> set encapsulation-protocol = ppp


```
admin> set ip-options source-if = sip0
admin> set telco-options call-type = off
admin> set ppp-options send-auth-mode = pap-ppp-auth
admin> set ppp-options recv-password = ascend
admin> set pppoe-options pppoe = yes
admin> write -f
```

When the PPPoE connection is established, the following command shows that the address assigned to the remote device is in the address space assigned to address pool 12:

```
admin> ifmgr -d | grep pppoe2
bif slot sif u m p ifname      host-name  remote-addr      local-addr
-----
044 2:01 003 *      wan44      pppoe2      194.194.2.1/32    220.220.220.112/32
```

Expanding the sample configuration to use a virtual router

The system generates default `slot-vr-config` profiles for the global router, but to use slot-based address assignment with a virtual router, you must create a new `slot-vr-config` profile.

For example, suppose the LIM connection in Figure 4-25 (page 4-95) must be partitioned in a secure virtual domain using a virtual router named `vr1`.

In addition to the configuration described in “Sample slot-based address configuration using the global router” on page 4-37, you need to complete the following steps:

- 1 Create a virtual router named `vr1` and configure address pools for that router. For example:

```
admin> new vrrouter
admin> set name = vr1
admin> set pool-summary = yes
admin> set pool-selection = slot-based
admin> set pool-base-address 4 = 195.195.1.1
admin> set pool-base-address 5 = 195.195.2.1
admin> set pool-base-address 6 = 195.195.3.1
admin> set assign-count 4 = 254
admin> set assign-count 5 = 254
admin> set assign-count 6 = 254
admin> write -f
```

- 2 Create a new `slot-vr-config` profile for the LIM slot. Specify the virtual router named `vr1` and a configured address pool. For example:

```
admin> new slot-vr-config
admin> set index physical-address shelf = shelf-1
admin> set index physical-address slot = slot-2
```

```
admin> set index vrouter = vr1
admin> set address-pool = 5
admin> write -f
```

3 Modify the PPPoE connection profile to use the virtual router named **vr1**. For example:

```
admin> read connection pppoe2
admin> set vrouter = vr1
admin> write -f
```

Configuring IP-ROUTE profiles

Any profile that specifies how to reach an IP device or subnet (such as an IP-INTERFACE, CONNECTION, or RADIUS user profile) specifies a static IP route to that destination. However, you can also configure static routes explicitly, to extend or fine-tune the routing table.

Overview of typical static route settings

You can define static routes in IP-ROUTE profiles or in RADIUS. For details about RADIUS pseudo-user and user profile route definitions, using the Framed-Route (22) or Ascend-Private-Route (104) attribute-value pair, see the *TAOS RADIUS Guide and Reference*.

Following are the local parameters (shown with default settings) for configuring a static route:

```
[in IP-ROUTE/""]  
name* = ""  
dest-address = 0.0.0.0/0  
gateway-address = 0.0.0.0  
metric = 8  
private-route = no  
active-route = yes
```

Parameter	Settings
name	Name of the profile (up to 31 characters).
dest-address	Destination IP address. The default value is 0.0.0.0, which represents the default route.
gateway-address	IP address of a next-hop router used to reach the specified destination. A next-hop router is directly connected to the same Ethernet segment, or is one hop away on a WAN link. When the Gigabit Ethernet interface is connected to a subnet and RIP is not enabled on the interface, the system must be informed about the gateway-address of other backbone routers that can route beyond the subnet.

Parameter	Settings
metric	RIP metric (0–15) for the route. Among routes with the same destination address, the higher the metric, the less likely that the system will choose the route.
private-route	Enables or disables including the route in RIP updates.
active-route	Enables or disables entering the route in the routing table. (Setting the parameter to <code>no</code> is a useful way to make a route temporarily inactive, so you can reinstate the route later.)


Offloading routing overhead to an external router

To offload routing overhead from the Stinger unit, you can define a default route to a router on the Gigabit Ethernet interface. A default route is a special-case static route that acts as a catch-all for packets for which the Stinger unit cannot find a route. A default route has the zero address as its destination and points to a specific gateway address. The system routes all packets with unknown destinations to the specified gateway. If no default route is defined, the system drops those packets.

The system creates an IP-ROUTE profile named `default`, but the profile is not valid until you specify a gateway address, so the route is not active until you assign an address and activate the route. You can create a default route by modifying the `default` profile, or by creating one or more IP-ROUTE profiles that specify a zero destination and a valid gateway address.

Figure 4-5 shows a router that resides on the same subnet as the Gigabit Ethernet IP interface. In this example, the system offloads part of its routing overhead by using a default route to the LAN router.

Figure 4-5. Default route to a local IP router




The following commands define a default route to the local router:

```
admin> read ip-route default
admin> set gateway-address = 1.1.1.200
admin> write -f
```

Creating a static route to a subnet

When RIP is turned off on an IP interface, the router cannot reach subnets beyond other routers on that interface unless it has a static route to the subnet. To enable access to subnets beyond the local segment, you must configure a static route. Figure 4-6 shows an example.

Figure 4-6. Static route to a subnet



The following commands configure a static route to the remote subnet:

```
admin> new ip-route subnet
admin> set dest-address = 135.1.2.3
admin> set gateway-address = 1.1.1.10
admin> write -f
```

Overview of routed subscriber connection features

The system creates a routing interface for local CONNECTION profiles when it starts up. For interfaces that use pool addresses or are defined in RADIUS user profiles, the system creates a routing interface when a session becomes active.

Subscriber connections can be PPPoA, PPPoE, IPoA, or BIR connections.

Source interface local addresses

You can define a Stinger local address to be used by multiple subscriber connections across DSL interfaces.



Note This feature is intended for use in CONNECTION or RADIUS profiles on Stinger DSL interfaces only.

The `source-if` setting provides several advantages over use of the `local-address` field as a means of assigning a connection-specific local address. For example, with `source-if` fewer IP addresses are consumed, because you can use the same local source address for many subscriber connections. In addition, the `source-if` address is used as the source address in most system-generated packets, which provides a security enhancement for sites that prefer not to reveal the configured `system-ip-addr` to packet analysis software.

Packets that use the specified source address

The system uses the assigned `source-if` address for all packets it generates except those in which upper protocol layers have already set the source IP address in a packet as, for example, in ICMP packets generated by the `ping -x` command. For those packets, the source address is set as it was in previous software versions to the

value of the `system-ip-addr` parameter in the IP-GLOBAL profile, or to the management interface IP address.

CPE client considerations

The subscriber connections can be routed IPoA, PPPoA, PPPoE, or BIR connections. For PPP connections, the system uses the assigned source address during IP NCP negotiation.

Multiple BIR connections can be configured with the same `local-address` setting. Support for that feature remains unchanged.

Soft IP interface requirement

With the `source-if` feature, for the system to bring up a connection to a CPE, it must have a configured soft IP interface for the `source-if` interface. You configure a soft IP interface by creating a new IP-INTERFACE profile with the default zero physical address and a nonzero logical item number. For example, the next commands create the logical soft interface and set its address to 10.0.0.254:

```
admin> new ip-interface
admin> set interface-address logical-item = 1
admin> set ip-address = 10.0.0.254
admin> write -f
```

To assign this source address to a subscriber connection, you specify the name of the interface in the `source-if` parameter. The name uses the following format:

`sipN`

The `N` is the `logical-item` value of the profile. For example, the default soft interface is `sip0`, and the soft interface created in the preceding commands is `sip1`. You can display interface names by using the `netstat -i` command. For example:

```
admin> netstat -i
```

Name	MTU	Net/Dest	Address	Ipkts	Ierr	Opkts	Oerr
ie0	1500	210.210.210.0/24	210.210.210.127	14044	0	5808	0
lo0	1500	127.0.0.1/32	127.0.0.1	5798	0	5798	0
rj0	1500	127.0.0.2/32	127.0.0.2	0	0	0	0
bh0	1500	127.0.0.3/32	127.0.0.3	0	0	0	0
wanabe	1500	127.0.0.3/32	127.0.0.3	0	0	0	0
local	65535	127.0.0.1/32	127.0.0.1	6079	0	6079	0
mcast	65535	224.0.0.0/4	224.0.0.0	0	0	0	0
tunnel0	1500	127.0.0.5/32	127.0.0.5	0	0	0	0
vr0_main	1500	127.0.0.4/32	127.0.0.4	0	0	0	0
sip0	65535	-	-	0	0	0	0
sip1	65535	10.0.0.254/32	10.0.0.254	0	0	0	0

Overview of configuration settings

Following are the relevant parameters, shown with default settings:

```

[in IP-INTERFACE/{ { any-shelf any-slot 0 } 0 }
ip-address = 0.0.0.0/0

[in IP-INTERFACE/{ { any-shelf any-slot 0 } 0 }:interface address
physical-address = { any-shelf any-slot 0 }
logical-item = 0

[in CONNECTION/"" :ip-options]
source-if = ""


```

Parameter	Setting
ip-address	Address to be used as the source interface address for multiple subscriber connections. If you change this address while active calls refer to it through the <code>source-if</code> field, all of the calls are terminated and then reestablished.
physical-address	For a soft IP interface, this field must be set to its default value of <code>{ any-shelf any-slot 0 }</code> .
logical-item	For a soft IP interface to be used with <code>source-if</code> , this field must be set to a nonzero number.
source-if	Name of the configured soft IP interface whose IP address will be the local address for this connection. The name uses the format <code>sipN</code>, where <code>N</code> is the <code>logical-item</code> value of the soft interface's IP-INTERFACE profile. You can display interface names by using the <code>netstat -i</code> command. If the corresponding soft IP interface does not exist, the call disconnects and the system logs a disconnect code. The <code>source-if</code> and <code>local-address</code> parameters cannot both be set in the same profile. If you change the setting of this parameter for a call that is already in the UP state, the call is terminated and then reestablished. For PPPoA/PPPoE calls, the client must reinitiate the PPP negotiation. The Ascend-IP-Source-If (309) attribute provides this functionality in a RADIUS user profile.

Sample configuration with a source interface address

Figure 4-7 shows a sample configuration with groups of subscriber connection sharing a `source-if` address. In this case, packets generated by the Stinger unit will not use the 2.2.2.2 `system-ip-addr` configured in the IP-GLOBAL profile. A different source address will be used for each of the CPE groups.

Figure 4-7. Sample configuration using two soft IP interface addresses



The following commands configure the two soft IP interface profiles:

```
admin> new ip-interface { { 0 0 0 } 1 }
admin> set ip-address = 10.0.0.254/24
admin> write -f
admin> new ip-interface { { 0 0 0 } 2 }
admin> set ip-address = 15.9.7.254/24
admin> write -f
```

The next command displays interface names. Notice that the last two lines of the command output specify the sip1 and sip2 interfaces, which were created by the preceding set of commands.

```
admin> netstat -i
```

Name	MTU	Net/Dest	Address	Ipkts	Ierr	Opkts	Oerr
ie0	1500	12.9.56.0/24	12.9.56.81	14044	0	5808	0
ie1	1500	210.210.210.0/24	210.210.210.127	14044	0	5808	0
lo0	1500	127.0.0.1/32	127.0.0.1	5798	0	5798	0
rj0	1500	127.0.0.2/32	127.0.0.2	0	0	0	0
bh0	1500	127.0.0.3/32	127.0.0.3	0	0	0	0
wanabe	1500	127.0.0.3/32	127.0.0.3	0	0	0	0
local	65535	127.0.0.1/32	127.0.0.1	6079	0	6079	0
mcast	65535	224.0.0.0/4	224.0.0.0	0	0	0	0
tunnel0	1500	127.0.0.5/32	127.0.0.5	0	0	0	0
vr0_main	1500	127.0.0.4/32	127.0.0.4	0	0	0	0
sip0	65535	-	-	0	0	0	0
sip1	65535	10.0.0.254/32	10.0.0.254	0	0	0	0
sip2	65535	15.9.7.254/32	15.9.7.254	0	0	0	0

The next set of commands configures CONNECTION profiles for the CPE units labeled pipe1 and pipe2 in Figure 4-7. These two CPE units share the same source-if address.

```
admin> new connection pipe1
admin> set active = yes
admin> set encapsulation-protocol = atm
```

```
admin> set ip-options remote-address = 10.0.0.1/24
admin> set ip-options source-if = sip1
admin> set atm-options vci = 37
admin> set atm-options nailed-group = 301
admin> write -f
admin> new connection pipe2
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 10.0.0.2/24
admin> set ip-options source-if = sip1
admin> set atm-options vci = 37
admin> set atm-options nailed-group = 302
admin> write -f
```

The next set of commands configures CONNECTION profiles for the CPE units labeled pipe3 and pipe4 in Figure 4-7:

```
admin> new connection pipe3
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 15.9.7.1/24
admin> set ip-options source-if = sip2
admin> set atm-options vci = 37
admin> set atm-options nailed-group = 303
admin> write -f
admin> new connection pipe4
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 15.9.7.2/24
admin> set ip-options source-if = sip2
admin> set atm-options vci = 37
admin> set atm-options nailed-group = 304
admin> write -f
```


Anti-spoofing protection for IPoA, BIR, PPPoA, and PPPoE connections

IP spoofing is a method of sending IP packets with a faked IP source address to gain unauthorized access to systems, usually to cause problems such as denial-of-service attacks. With this software version, a simple anti-spoofing method is provided to protect the system from IP spoofing attempts across PPPoA, PPPoE, IPoA, and BIR connections.

Two levels of anti-spoofing protection are provided.

- **Level 1: Source IP address checks**

At level 1, the system checks the source address of a packet against the remote address of the CONNECTION or RADIUS user profile, and drops packets whose source address doesn't match. The source IP address check is recommended for a connection to a CPE that connects to a single subnet, or for a restrictive access that allows only packets from the CPE itself or a host on its subnet across the WAN interface.

The profile's address can be specific to a single host or a LAN, and the address can be statically specified in the profile or dynamically allocated via DHCP or address pools. However, source address checks cannot be used with DHCP over BIR connections.



Note Using source IP address checks or reverse path checks on a DHCP over BIR connection prevents the virtual interface from coming up.

- **Level 2: Reverse path checks**

At level 2, the system performs the source address check of level 1, but also verifies reachability of the source address across the same interface on which the packet was received. If a route to the packet's source address is not found, the system drops the packet. The reverse path check is recommended for connections to a router that services different networks, or for a less restrictive access that allows addresses beyond the CPE subnet to be accepted on the WAN interface.

Routes to source addresses can be configured explicitly through static routes or learned dynamically through any of the available routing protocols.

Anti-spoofing works on DSL and trunk interfaces, but does not work on the Gigabit Ethernet interface.

Overview of anti-spoofing settings

For PPP calls (PPPoE or PPPoA), anti-spoofing settings must be configured in the PPP CONNECTION profile (not in the ATM circuit CONNECTION profile between the LIM interface and the internal ATM interface of the LIM). For IPoA or BIR connections, anti-spoofing settings must be configured in their respective CONNECTION profiles.

Following are the relevant parameters, shown with default values:


```
[in CONNECTION/"" :ip-options]
source-ip-check = no
reverse-path-check = no
```

Parameter	RADIUS attribute	Setting
source-ip-check	Ascend-Source-IP-Check (96)	Enables or disables level 1 anti-spoofing protection. When it is enabled, the system compares the source address in packets received on the WAN interface to the profile's remote address, and drops nonmatching packets. By default, this protection is disabled. This setting and the reverse-path-check setting are mutually exclusive. Only one of these fields can be enabled in a profile. Use this setting to restrict traffic on the interface to packets from the profile's remote host or subnet.
reverse-path-check	Ascend-Reverse-Path-Check (310)	Enables or disables level 2 anti-spoofing protection. When it is enabled, the system checks the routing table to verify that it can reach the packet's source address across the same interface on which it received the packet, and drops nonmatching packets. By default, this protection is disabled. This setting and the source-ip-check setting are mutually exclusive. Only one of these fields can be enabled in a profile. Use this setting to allow all packets from the profile's remote host or subnet, as well as packets received from networks that are reachable through that interface.

Sample anti-spoofing configuration

With the sample network shown in Figure 4-8, the CPE must have a CONNECTION profile in the Stinger system. The profile can enable `source-ip-check` (level 1 protection) or `reverse-path-check` (level 2 protection) on a PPPoA call.

Figure 4-8. Sample network with two levels of anti-spoofing protection



- If `source-ip-check` is enabled for the subscriber connection, packets from 2.2.2.x/24 subnet will be accepted, but packets from 3.3.3.x network will be dropped.
- If `reverse-path-check` is enabled for the subscriber connection and a static route for 3.3.3.x is configured (or a routing protocol is enabled on the WAN interface to

the CPE), packets from both the 2.2.2.x/24 subnet and the 3.3.3.x network will be accepted.

Before you begin, enable an DSL interface and configure an ATM circuit from the interface to the LIM's ATM internal interface. For details about the ATM circuit configuration for PPPoA, see "Configuring an ATM circuit to terminate PPPoA" on page 4-90.

The following commands create a PPPoA CONNECTION profile for level 1 anti-spoofing. Only packets from the CPE and hosts on its subnet will be accepted.

```
admin> new connection pppoa-1
admin> set active = yes
admin> set encapsulation-protocol = ppp
admin> set telco-options call-type = off
admin> set ppp-options recv-password = localpw
admin> set ip-options remote-address = 2.2.2.4/24
admin> set ip-options source-ip-check = yes
admin> write -f
```

Configuring IPoA subscriber connections

The CPE devices described in this section are IP-capable DSL devices that transmit IP over ATM (IPoA). IPoA is a type of RFC 2684 traffic.

The CPE devices require a terminating PVC to the Stinger unit. A terminating PVC is not switched through the system, it terminates on the IP control module and its data stream is passed up to the IP router for further handling.

The profile for a terminating connection must specify the IP address of the far-end router, and it can set a number of other routing-related values. The profile must also specify the ATM characteristics of the connection (for example, a VPI and VCI assignment and a nailed group representing the interface to use). The *Stinger ATM Configuration Guide* describes the ATM aspects of the configuration in detail.

Typical CONNECTION atm-options settings for terminating PVCs

For a discussion of ATM settings and quality of service (QoS) contracts, see the *Stinger ATM Configuration Guide*. Following are the ATM-related parameters, shown with default settings, for ATM terminating PVCs:

```
[in CONNECTION/"" ]
station = ""
active = no
encapsulation-protocol = atm-circuit

[in CONNECTION/"" :atm-options]
atm1483type = aal5-llc
vpi = 0
vci = 35
nailed-group = 1
```

Parameter	RADIUS attribute	Setting
station	User-Name (1)	Name of the far-end device.
active	N/A	Enables or disables the profile.
encapsulation-protocol	Framed-Protocol (7)	Encapsulation protocol to use for the connection. Must specify ATM for terminating PVCs.
atm1483type	Framed-Protocol (7)	Stinger systems support the two encapsulation methods for carrying routed PDUs in the payload field of ATM adaptation layer (AAL) type 5, which are defined in RFC 2684, <i>Multiprotocol Encapsulation over ATM Adaptation Layer 5</i>. For details about how to use this setting, see “PPPoA and PPPoE encapsulation types in the underlying ATM circuit” on page 4-89.
vpi	Ascend-ATM-Vpi (94)	VPI value for the PVC. For a discussion of valid values, see the <i>Stinger ATM Configuration Guide</i> .
vci	Ascend-ATM-Vci (95)	VCI value for the PVC. For a discussion of valid values, see the <i>Stinger ATM Configuration Guide</i> .
nailed-group	Ascend-ATM-Group (64)	Nailed-group number of the interface used by the connection. You can obtain the nailed-group assigned to any interface by using the <code>which -n</code> command.

Typical CONNECTION ip-options settings for terminating PVCs

For information about enabling IP multicast forwarding on client or remote MBONE interfaces, see Chapter 9, “IP Multicast Configuration.”

Following are the IP options (shown with default settings) for configuring an IP routed RFC 2684 connection to a DSL CPE:

```
[in CONNECTION/"" :ip-options]
ip-routing-enabled = yes
vj-header-prediction = no
remote-address = 0.0.0.0/0
local-address = 0.0.0.0/0
routing-metric = 1
private-route = no
rip = routing-off
```

Parameter	RADIUS attribute	Setting
ip-routing-enabled	Ascend-Route-IP (228)	Enables or disables IP routing on the interface.
vj-header-prediction	Framed-Compression (13)	Enables or disables Van Jacobson prediction for TCP packets on incoming calls using encapsulation protocols that support Van Jacobson compression.


Parameter	RADIUS attribute	Setting
remote-address	Framed-IP-Address (8) Framed-IP-Netmask (9)	IP address of the remote CPE device.
local-address	Ascend-PPP-Address (253) Ascend-IF-Netmask (153)	Local IP address of a numbered interface connection. For a more flexible alternative to this setting, see “Source interface local addresses” on page 4-42.
routing-metric	Ascend-Metric (225)	RIP metric (1–15) for the specified route. If preference values are equal, the higher the metric, the less likely that the router will use the route.
private-route	Ascend-Route-Preference (126)	Include or exclude this route in RIP updates.
rip	Framed-Route (22)	Enables or disables RIP.

For details about parameters, see the *Stinger Reference*. For details about the attribute-value pairs used to configure IP options in RADIUS profiles, see the *TAOS RADIUS Guide and Reference*.

Sample RFC 2684 (IPoA) terminating PVC

Figure 4-9 shows a CPE router connection using IPoA.

Figure 4-9. Router-to-router IP connection



The default settings for the `ip-options` subprofile enable IP routing and Van Jacobson header compression and turn RIP off. Those settings are typically appropriate for a DSL interface, but they are not required. The following example shows configuration of a `CONNECTION` profile for the DSL CPE router in Figure 4-9:

```
admin> read connection router-1
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 10.7.8.200/30
admin> set atm-options vpi = 8
admin> set atm-options vci = 100
```


```
admin> set atm-options nailed-group = 201
admin> write -f
permconn-st-1 Password = "ascend"
  Service-Type = Outbound,
  Framed-Protocol = ATM-1483,
  User-Name = "router-1",
  Framed-IP-Address = 10.7.8.200,
  Framed-IP-Netmask = 255.255.255.252,
  Ascend-ATM-Group = 201,
  Ascend-Route-IP = Route-IP-Yes,
  Ascend-ATM-Vci = 100
```

Example of using a local-address setting for a numbered interface

A numbered-interface configuration assigns each side of the connection a unique address that applies only to that connection. Figure 4-10 shows a numbered-interface connection. The Stinger unit's real, physical Ethernet interface has the IP address 2.2.2.2. The other two addresses represent the local and remote addresses of the numbered-interface connection.

See "Source interface local addresses" on page 4-42 for information about local addresses that can be used by multiple subscriber connections across DSL interfaces.

Figure 4-10. A numbered-interface connection



Some applications such as SNMP use the local-address value internally to keep track of the circuit. The local-address value must be unique to the connection and to the network.



Note Do not assign a local address that belongs to one of the Stinger unit's real physical LAN interfaces. Doing so causes routing problems.

The following set of commands specifies a CONNECTION profile for the numbered interface:

```
admin> new connection numbered
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 3.3.3.3/30
admin> set ip-options local-address = 7.7.7.7/30
admin> set atm-options vci = 36
```

```
admin> set atm-options nailed-group = 211
admin> write -f
```

Following is a comparable RADIUS profile:


```
permconn-st-2 Password = "ascend"
  Service-Type = Framed-User,
  Framed-Protocol = ATM-1483,
  User-Name = "numbered",
  Ascend-ATM-Group = 211,
  Ascend-Route-IP = Route-IP-Yes,
  Ascend-ATM-Vpi = 0,
  Ascend-ATM-Vci = 36,
  Framed-IP-Address = 3.3.3.3,
  Framed-IP-Netmask = 255.255.255.252,
  Ascend-PPP-Addr = 7.7.7.7,
  Ascend-IF-Netmask = 255.255.255.252
```

In this example, the interface is assigned a 30-bit subnet, so four bit combinations are available for host assignments. Of the four possible host addresses, the one that is evenly divisible by 4 is the network or base address (the address that specifies zeros in the host bits). This address is added to the routing table. The other host addresses are assigned a /32 subnet mask and added as host routes. You can suppress advertisement of the host routes associated with a numbered interface by using the `suppress-host-routes` parameter, which is described in the *Stinger Reference*.

Example of routing a terminated PVC across Gigabit Ethernet

You can forward IPoA from DSL subscribers onto the Gigabit Ethernet IP interface to be further routed to a specific IP destination such as an Internet service provider (ISP), as shown in Figure 4-11.

Figure 4-11. Forwarding terminating PVCs on the Gigabit Ethernet interface



This configuration requires a terminating PVC for each DSL subscriber. In this example, the Stinger does not maintain a large routing table itself. It uses a static route configuration to forward IP traffic across Gigabit Ethernet to another router, which routes the traffic on toward the ISP point of presence.

The following commands create a CONNECTION profile for each of the DSL subscribers in Figure 4-11:

```
admin> new connection user-1
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 2.2.2.2/29
```

```
admin> which -n { 1 2 1 }
Nailed group corresponding to port { shelf-1 slot-2 1 } is 51
admin> set atm-options nailed-group = 51
admin> write -f
admin> new connection user-2
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 3.3.3.3/29
admin> which -n { 1 2 2 }
Nailed group corresponding to port { shelf-1 slot-2 2 } is 52
admin> set atm-options nailed-group = 52
admin> write -f
```

The following command displays the IP-INTERFACE profile for the Gigabit Ethernet interface, which shows that the address has been specified and RIP is not enabled:

```
admin> get ip-interface { { 1 8 2 } 0 }
[in IP-INTERFACE/{ { shelf-1 first-control-module 2 } 0 }]
interface-address* = { { shelf-1 first-control-module 2 } 0 }
ip-address = 1.1.1.1/28
proxy-mode = Off
rip-mode = routing-off
route-filter = ""
rip2-use-multicast = yes
ospf = { no 0.0.0.0 normal 10 40 5 simple ***** 0 1 16777215 type-1 c+
multicast-allowed = no
igmp-options = { 2 125 100 10 2 }
multicast-rate-limit = 100
multicast-group-leave-delay = 0
multicast-group-leave-delay-msec = 0
multicast-service-profile = ""
multicast-max-groups = 0
directed-broadcast-allowed = yes
vrouter = ""
management-only-interface = no
vlan-enabled = no
vlan-id = 0
```

The following set of commands configures a static route to the ISP's destination address, specifying a next-hop router on the Gigabit Ethernet interface:


```
admin> read ip-route isp-dest
admin> set dest-address = 7.7.7.7/24
admin> set gateway-address = 1.1.1.70
admin> set active-route = yes
admin> write -f
```

With this example configuration, when packets destined for 7.7.7.7/24 are received on the terminating PVCs, the IP DSLAM consults its own routing table and forwards the packets onto a Gigabit Ethernet interface to the next-hop router specified as the gateway-address.

Example of using IP routing to aggregate PVCs onto a trunk VC

You can use IP routing to aggregate many IPoA connections from DSL subscribers onto a single virtual circuit to a specific IP destination such as an ISP. Instead of configuring an ATM circuit for each subscriber, you use PVCs that terminate on the IP DSLAM and use IP routing to direct the traffic out on a terminating PVC to the ISP. This greatly simplifies provisioning new DSL subscribers that route to the same ISP.

Figure 4-12. Aggregating PVCs onto a single virtual circuit using IP routing



The following commands create a CONNECTION profile for each of the DSL subscribers in Figure 4-12:

```
admin> new connection user-1
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 2.2.2.2/29
admin> set atm-options vpi = 0
admin> set atm-options vci = 36
admin> set atm-options nailed-group = 51
admin> write -f
admin> new connection user-2
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 3.3.3.3/29
admin> set atm-options vpi = 0
admin> set atm-options vci = 37
admin> set atm-options nailed-group = 52
admin> write -f
```

The next command configures the PVC to the ISP:

```
admin> new connection isp
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 7.7.7.7/24
admin> set atm-options vpi = 0
admin> set atm-options vci = 35
admin> which -n { 1 17 2}
```


```
Nailed group corresponding to port { shelf-1 trunk-module-1 2 } is 802
admin> set atm-options nailed-group = 802
admin> write -f
```

This CONNECTION profile creates a static route to the ISP's destination address across the trunk interface. When packets destined for 7.7.7.7/24 are received on the terminating PVCs, the IP DSLAM consults its own routing table and forwards the packets onto the trunk interface to the ISP.

Configuring BIR subscriber connections

With bridged IP routing (BIR), a Stinger IP DSLAM can establish an IP routed connection to an IP host through a customer premises equipment (CPE) bridge device. A BIR connection can use a line interface module (LIM) or trunk interface. A sample setup with a BIR interface on a LIM port is shown in Figure 4-13.

Figure 4-13. BIR interface on a LIM port



On the BIR interface, the system receives IP packets encapsulated in bridged frames. The IP control module decapsulates the packets and passes them up the protocol stack to the IP router. To the IP host, the session appears to be an ordinary IP connection.

BIR configurations require the use of numbered interfaces, which assign both the remote and local side of the connection a unique IP address. The remote address can specify a subnet or an individual remote IP host. Typically, the local address for the Stinger unit is a unique address on the remote subnet. For details about numbered interfaces, see “Example of using a local-address setting for a numbered interface” on page 4-52. For related information, see “Source interface local addresses” on page 4-42.

Overview of CONNECTION bir-options and ip-options settings

In addition to the many possible IP routing parameters in CONNECTION and RADIUS profiles, described in “Configuring IPoA subscriber connections” on page 4-49, the following parameters apply to BIR interfaces. The parameters are shown with default settings.

```
[in CONNECTION/"" :bir-options]
enable = no
proxy-arp = no
multiple-ip-address = { no [ { 0.0.0.0/0 0.0.0.0/0 } { 0.0.0.0/0 0.0.0.0/0 +
[in CONNECTION/"" :ip-options]
ip-routing-enabled = yes
remote-address = 0.0.0.0/0
```

local-address = 0.0.0.0/0


Parameter	RADIUS attribute	Setting
enable	Ascend-BIR-Enable (70)	Enables or disables BIR on this interface.
proxy-arp	Ascend-BIR-Proxy (71)	Enables or disables proxy Address Resolution Protocol (ARP), which causes the Stinger IP DSLAM to respond as proxy for ARP requests from local hosts for remote hosts on the far end of the link.
multiple-ip-address NA		The multiple-ip-address subprofile contains a four-element array, with each element representing one pair of remote and local IP addresses for the static configuration of up to four WAN virtual interfaces. For details, see “Configuring multiple WAN virtual IP interfaces on a BIR connection” on page 4-62.
ip-routing-enabled	Ascend-Route-IP (228)	Enables or disables IP routing on the interface.
remote-address	Framed-IP-Address (8) Framed-IP-Netmask (9)	IP address of the remote device, which can include a subnet specification. If the address does not include a subnet mask, the router assumes the default subnet mask based on address class.
local-address	Ascend-PPP-Address (253) Ascend-IF-Netmask (153)	IP address assigned to the local side of a numbered-interface connection. This is a requirement for BIR interfaces. For a more flexible alternative setting, see “Source interface local addresses” on page 4-42.

Sample subnet (BIR/24) configuration

A BIR subnet configuration specifies a remote subnet address, and can be used to transmit bridged data to multiple IP hosts on that subnet.

When the Stinger IP DSLAM receives a packet destined for a BIR subnet interface, it examines the network bits of the destination address and forwards the packet to the related CPE. For example, Figure 4-14 shows two bridging CPE devices connected to an IP class C subnet. With this example, if the IP DSLAM receives a packet addressed to 2.2.2.200 or 2.2.2.128, it examines only the first 24 bits of the address, and forwards the packets to the bridge CPE.

Figure 4-14. BIR subnet configuration on LIM interface



The following commands configure a BIR subnet interface through the DSL CPE bridge in Figure 4-14:

```
admin> new connection bir-1
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 2.2.2.0/24
admin> set ip-options local-address = 2.2.2.1/32
admin> set bir-options enable = yes
admin> set bir-options proxy-arp = yes
admin> set atm-options atm1483type = aal5-llc
admin> set atm-options vci = 101
admin> set atm-connect-options atm1483type = aal5-llc
admin> which -n { 1 2 1 }
Nailed group corresponding to port { shelf-1 slot-2 1 } is 51
admin> set atm-options nailed-group = 51
admin> write -f
```


Following is a comparable definition in a RADIUS profile:

```
permconn-cpe-1 Password = "ascend"
    Service-Type = Outbound,
    Framed-Protocol = ATM-1483,
    User-Name = "bir-1",
    Ascend-Route-IP = Route-IP-Yes,
    Framed-IP-Address = 2.2.2.0,
    Framed-IP-Netmask = 255.255.255.0,
    Ascend-PPP-Addr = 2.2.2.1,
    Ascend-IF-Netmask = 255.255.255.255,
    Ascend-ATM-Group = 51,
    Ascend-ATM-Vci = 101,
    Ascend-BIR-Enable = BIR-Enable-Yes,
    Ascend-BIR-Proxy = BIR-Proxy-Yes
```

Sample host route (BIR/32) configurations

When a Stinger IP DSLAM receives a packet to a BIR/32 interface, it examines the full 32 bits of the destination address and forwards the packet to the related CPE. Figure 4-15 shows two bridging DSL CPE devices, each supporting one host. In this example, the IP hosts have addresses on the same IP network, but that is not a requirement.

Figure 4-15. BIR/32 configurations



In Figure 4-15, the local-address value is the same for both BIR interfaces. This is recommended for host routes to the same IP network because it simplifies configuration of the remote hosts, all of which can point to the same local address as the gateway.

The following commands configure a BIR/32 interface through the CPE labeled *Bridge-5* in Figure 4-15:

```
admin> new connection bir-5
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 4.4.4.5/32
admin> set ip-options local-address = 4.4.4.1/32
admin> set atm-options atm1483type = aal5-llc
admin> set atm-options vci = 111
admin> set bir-options enable = yes
admin> set atm-connect-options atm1483type = aal5-llc
admin> which -n { 1 2 5 }
Nailed group corresponding to port { shelf-1 slot-2 5 } is 55
admin> set atm-options nailed-group = 55
admin> write -f
```

The following commands modify the CONNECTION profile immediately above to configure a BIR/32 interface through the CPE labeled *Bridge-6*:

```
admin> set station = bir-6
(New index value; will save as new profile CONNECTION/bir-6.)
admin> set ip-options remote-address = 4.4.4.6/32
admin> set atm-options vci = 112
admin> which -n { 1 2 6 }
Nailed group corresponding to port { shelf-1 slot-2 6 } is 56
admin> set atm-options nailed-group = 56
admin> write -f
```

Following are comparable definitions in RADIUS profiles:

```
permconn-cpe-5 Password = "ascend"
    Service-Type = Outbound,
    Framed-Protocol = ATM-1483,
    User-Name = "bir-5",
```

```
Ascend-Route-IP = Route-IP-Yes,  
Framed-IP-Address = 4.4.4.5,  
Framed-IP-Netmask = 255.255.255.255,  
Ascend-PPP-Addr = 4.4.4.1,  
Ascend-IF-Netmask = 255.255.255.255,  
Ascend-ATM-Group = 55,  
Ascend-ATM-Vci = 111,  
Ascend-BIR-Enable = BIR-Enable-Yes  
permconn-cpe-6 Password = "ascend"  
Service-Type = Outbound,  
Framed-Protocol = ATM-1483,  
User-Name = "bir-6",  
Ascend-Route-IP = Route-IP-Yes,  
Framed-IP-Address = 4.4.4.6,  
Framed-IP-Netmask = 255.255.255.255,  
Ascend-PPP-Addr = 4.4.4.1,  
Ascend-IF-Netmask = 255.255.255.255,  
Ascend-ATM-Group = 56,  
Ascend-ATM-Vci = 112,  
Ascend-BIR-Enable = BIR-Enable-Yes
```

Sample BIR connection on a VDSL port

Very high speed DSL (VDSL) ports use Ethernet encapsulation and QoS on the drop rather than ATM. This example shows configuration of a BIR/32 interface on a VDSL port using Ethernet encapsulation and QoS.

The following commands configure an ETHER-QOS profile with priority 3 (the highest priority) and a committed information rate (CIR) of 512Kbps. This is equivalent to UBR service in ATM QoS.

```
admin> new ether-qos sats-dn  
admin> set priority = 3  
admin> set committed-information-rate = 512  
admin> write -f
```

The following commands configure a BIR CONNECTION profile that applies the sats-dn ETHER-QOS profile in the downstream direction on VDSL interface 701:

```
admin> new connection bir-15-1  
admin> set active = yes  
admin> set encapsulation-protocol = ethernet  
admin> set ip-options remote-address = 64.1.1.24/32  
admin> set ip-options multicast-allowed = yes  
admin> set ip-options igmp-options robust-count = 10  
admin> set ip-options igmp-options query-interval = 1024  
admin> set bir-options enable = yes  
admin> set bir-options proxy-arp = yes  
admin> set ether-options nailed-group = 701
```

```
admin> set ether-options usr-up-stream-contract = default
admin> set ether-options usr-dn-stream-contract = sats-dn
admin> write -f
```


For background information about the ETHER-QOS profile and Ethernet options in CONNECTION profiles, see the *Stinger VDSL Line Interface Module (LIM) Guide*.

Sample use of filters with BIR connections

You can apply an IP filter to restrict outbound packets on a BIR interface. However, IP filters are not applied to the inbound packet stream on a BIR interface. For details about defining IP filters, see Chapter 11, “Filter Configuration.”

Figure 4-16 shows a sample BIR interface to a subnet that supports two IP hosts.

Figure 4-16. Bidirectional filtering on a BIR interface



The filter defined in this example is applied to the BIR interface. The input filter rules affect packets received on the BIR interface, and output filter rules affect packets destined for the user-side subnet.

The first input filter rule shown below specifies that if the destination IP address in a packet is 2.2.2.0/24, the protocol is 17 (UDP), and the source UDP port is less than 50, the packet is discarded. So, packets that match this rule will not reach the server at 2.2.2.36. The second input filter is an explicit default rule that forwards all other IP packets received on the BIR interface.

```
admin> new filter udp-filter
admin> set input-filters 1 valid-entry = yes
admin> set input-filters 1 forward = no
admin> set input-filters 1 Type = ip-filter
admin> set input-filters 1 ip-filter protocol = 17
admin> set input-filters 1 ip-filter dest-address-mask = 255.255.255.0
admin> set input-filters 1 ip-filter dest-address = 2.2.2.36
admin> set input-filters 1 ip-filter Src-Port-Cmp = less
admin> set input-filters 1 ip-filter source-port = 50
admin> set input-filters 2 valid-entry = yes
admin> set input-filters 2 forward = yes
admin> set input-filters 2 Type = ip-filter
```

The first output filter rule shown below specifies that if the source IP address in a packet is 2.2.2.36/24, the protocol is 17, and the source UDP port is less than 50, the packet is discarded. So, packets that match this rule will not reach the IP hosts across the BIR interface. The second output filter is an explicit default rule that forwards all other IP packets destined for the remote subnet through the BIR interface.

```
admin> set output-filters 1 valid-entry = yes
admin> set output-filters 1 forward = no
admin> set output-filters 1 Type = ip-filter
admin> set output-filters 1 ip-filter source-address-mask = 255.255.255.0
admin> set output-filters 1 ip-filter source-address = 2.2.2.36
admin> set output-filters 1 ip-filter protocol = 17
admin> set output-filters 1 ip-filter Src-Port-Cmp = less
admin> set output-filters 1 ip-filter source-port = 50
admin> set output-filters 2 valid-entry = yes
admin> set output-filters 2 forward = yes
admin> set output-filters 2 Type = ip-filter
admin> write -f
```

The following commands create a BIR profile to the bridge CPE in Figure 4-16, and apply the sample filter:

```
admin> new connection bir-1-1
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 1.1.1.0/24
admin> set ip-options local-address = 1.1.1.1/24
admin> set session-options data-filter = udp-filter
admin> set bir-options enable = yes
admin> set atm-options nailed-group = 101
admin> write -f
```

Configuring multiple WAN virtual IP interfaces on a BIR connection

You can statically configure up to four remote and local IP address pairs on a BIR connection. Each address pair is a WAN virtual interface.

If a filter is applied to the BIR CONNECTION profile, the filter is applied to the main BIR connection and all the WAN virtual interfaces inherit the filter rules.

The WAN virtual interfaces can be configured or reconfigured at runtime, without affecting the main connection or other WAN virtual interfaces on the connection. However, if an existing BIR connection has an address specified in the `ip-options:remote-address` field, and you reconfigure the connection to support multiple IP addresses instead, you must disable and reactivate the connection for the new settings to take effect.

Overview of multiple IP address settings

Following are the parameters, shown with default settings, for configuring multiple static addresses:


```
[in CONNECTION/""bir-options:multiple-ip-address]
enable = no
ip-address-option = [ { 0.0.0.0/0 0.0.0.0/0 } { 0.0.0.0/0 0.0.0.0/0 } { 0.+
[in CONNECTION/""bir-options:multiple-ip-address:ip-address-option]
ip-address-option[1] = { 0.0.0.0/0 0.0.0.0/0 }
ip-address-option[2] = { 0.0.0.0/0 0.0.0.0/0 }
ip-address-option[3] = { 0.0.0.0/0 0.0.0.0/0 }
ip-address-option[4] = { 0.0.0.0/0 0.0.0.0/0 }
[in CONNECTION/""bir-options:multiple-ip-address:ip-address-option[1]]
remote-address = 0.0.0.0/0
local-address = 0.0.0.0/
```

Parameter	Setting
enable	Enables or disables static configuration of multiple IP addresses for this BIR connection. Changing the value of this setting activates or deactivates all correctly configured WAN virtual interfaces configured in the <code>multiple-ip-address</code> subprofiles. If enabled, at least one <code>remote-address</code> field must be configured in one of the <code>multiple-ip-address</code> subprofiles. This parameter cannot be set to <code>yes</code> if the <code>bootp-relay-options enable</code> setting is set to <code>yes</code>, and vice versa.
ip-address-option	An array of four elements, with each element representing a remote and local address pair for one WAN virtual IP interface on this BIR connection.
remote-address	The remote address of a WAN virtual interface. This address must be unique across all Stinger configuration settings system wide. The <code>multiple-ip-address enable</code> setting must be set to <code>yes</code> to configure this parameter.
local-address	The local address, if required, of the WAN virtual interface. The IP address of the local virtual interface is obtained either from this setting, or if this address is null, from the <code>connection ip-options source-if</code> setting, or if that address is null, from the <code>connection ip-options local-address</code> setting. If no local address is available, the WAN virtual interface will not come up.

Sample configuration of multiple addresses on a BIR connections

The sample configuration in this section use the basic network setup and IP addresses shown in Figure 4-17:

Figure 4-17. Sample BIR connection with four static IP addresses



To configure the BIR connection, create a CONNECTION profile for the remote bridge CPE. Disable BOOTP-relay, and configure the remote and local address pairs statically. For example:

```
admin> read connection
admin> set station = bir-cpe-1
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options bootp-relay-options enable = no
admin> set atm-options vci = 36
admin> set atm-options nailed-group = 51
admin> set bir-options enable = yes
admin> set bir-options multiple-ip-address enable = yes
admin> set bir-options multiple ip-addr 1 remote-address = 1.1.1.5/8
admin> set bir-options multiple ip-addr 1 local-address = 1.1.1.2/8
admin> set bir-options multiple ip-addr 2 remote-address = 2.2.2.1/8
admin> set bir-options multiple ip-addr 2 local-address = 2.2.2.2/8
admin> set bir-options multiple ip-addr 3 remote-address = 3.3.3.3/8
admin> set bir-options multiple ip-addr 3 local-address = 3.3.3.2/8
admin> set bir-options multiple ip-addr 4 remote-address = 4.4.4.10/8
admin> set bir-options multiple ip-addr 4 local-address = 4.4.4.4/8
admin> write -f
```

When you write the profile, the main BIR connection to the bridge CPE is established first. The system generates a log message such as the following:

```
LOG info, Shelf 1, Slot 2, Time: 03:24:44--
[1/2/1/0] LAN session up: <bir-cpe-1> [MBID 11] [bir-cpe-1]
```

The system then creates WAN virtual interfaces for each of the configured valid addresses. For example, the following `ifmgr` command, which is available only in the debug environment, shows the initial connection and four WAN virtual interfaces:

```
super> ifmgr -d
bif slot sif u m p ifname      host-name  remote-addr      local-addr
-----
...
069 1:02 067 *   p wan69      bir-cpe-1  0.0.0.0/32       0.0.0.0/32
074 1:02 068 *   p wan69-1    bir-cpe-1  1.1.1.5/32       1.1.1.2/32
075 1:02 069 *   p wan69-2    bir-cpe-1  2.2.2.1/32       2.2.2.2/32
076 1:02 070 *   p wan69-3    bir-cpe-1  3.3.3.3/32       3.3.3.2/32
077 1:02 071 *   p wan69-4    bir-cpe-1  4.4.4.10/32      4.4.4.4/32
<end>
```

The WAN virtual interfaces can be killed by using the `ifmgr -k` command.

For details about WAN virtual interfaces configured via DHCP, see “WAN virtual interfaces on DHCP-configured BIR connections” on page 4-66.

Configuring DHCP relay for IPoA and BIR connections

With DHCP relay, the system transfers messages between a client requesting its IPoA or BIR configuration, and a DHCP server. To enable DHCP requests on BIR connections, the system creates one virtual interface for each remote DHCP client.

RFC compliance and caveats

Stinger support for DHCP relay with option 82 is compliant with the following RFCs, with the exception noted for RFC 1542:

- RFC 951, *Bootstrap Protocol (BOOTP)*
- RFC 1542, *Clarifications and Extensions for the Bootstrap Protocol*
- RFC 2131, *The Dynamic Host Configuration Protocol (DHCP)*
- RFC 2132, *DHCP Options and BOOTP Vendor Information Extensions*
- RFC 3046, *DHCP Relay Agent Information Option*



Note When WAN virtual interfaces are created on a BIR connection, the system is forced to adopt behavior that is noncompliant with section 4.1 of RFC 1542, which indicates that if a relay agent receives a DHCP packet for which the destination IP address is not one of the interfaces of the relay-agent, the packet must be forwarded without any changes. Instead, the Stinger catches the packet and adds the relay-agent information to it.

DHCP option 82

DHCP option 82, the relay agent Information option, is used to associate a unique identifier with a broadband device. The identifier can be associated with the virtual circuit to a remote device, or with the remote router itself. DHCP servers that recognize this option can use the option 82 identifier to enforce conditions on address or configuration access.

The identifier typically consists of the specified IP address concatenated to the connection name on which the DHCP request packet was received. The connection name is used by the relay agent to route back the packet.

WAN virtual interfaces on DHCP-configured BIR connections

To configure WAN virtual interfaces via DHCP for a BIR connection, all of the following conditions must be true:

- The system receives a DHCP reply packet containing option 82 from the DHCP server.
- The IP-GLOBAL profile `allow-multiple-interface-creation` field is set to `yes`.
- The BIR CONNECTION profile has DHCP relay enabled.
- The number of the created virtual interface is less than the configured one.
- If the `message-type` option is present in the DHCP packet, it must be set to DHCP-ACK. If the option is not present, the system assumes DHCP-ACK.

Direct routes associated with the virtual interface are set up during the interface creation. These dynamically created interfaces will be typed `virtual`, like virtual interfaces on an Ethernet port.

IP address assignments on DHCP-configured WAN virtual interfaces

Virtual interface addresses have the following characteristics:

- The IP address of the WAN virtual interface is obtained from the CONNECTION profile if the `local-address` setting is not null. Otherwise, it is extracted from option 3 (the router option) of the DHCP packet. The virtual interface is created with a /32 route.
- The destination IP address of a WAN virtual interface is obtained from the `yiaddr` field of the DHCP reply message.
- For the DHCP client obtaining its configuration from the DHCP server, the netmask is obtained from option 1 (the netmask option) if present in the DHCP reply packet. If option 1 is not provided, the netmask will be the default mask associated with the IP address class.
- For the DHCP client obtaining its configuration from the DHCP server, the broadcast address is obtained from option 28 (the broadcast option) if present in the DHCP reply packet. If option 28 is not provided, the broadcast address will be calculated from the destination IP address and the netmask. The host part of the destination address will have all bits set to 1.

DHCP IP address lease time for WAN virtual interfaces

The address for a WAN virtual IP interface is leased via DHCP for a number of seconds. The lease time is extracted from the DHCP packet, option 51, and can be extended for an interface by setting the `additional-lease-delay` parameter. The maximum lease time cannot exceed 42,949,672 seconds, which translates to 11930 hours (497 days). This limit is set because Stinger ticks are calibrated to 10 ms with a maximum value of a 32 bits/100 = 42949672. The Stinger interprets a lease time of 0 as an infinite lease.

If a WAN virtual interface is up and active for a given period, and a DHCP-ACK is received during its active period, then the lease time is updated by the new value. The interface then remains up for the new received lease time.

When WAN virtual interfaces are activated

The WAN virtual interfaces are activated if one of the following conditions occurs. Activating a virtual interface implies creation of associated routes.

- The interface has just been created.
- The interface was inactive and a DHCP reply has been received for the same host, same IP address.
- All associated virtual interfaces are activated when the main interface was inactive and becomes active. If the leased time has expired for an associated virtual interface, that virtual interface is deleted.

For details about how the system creates routes for WAN virtual interfaces, see “How routes are created for WAN virtual interfaces” on page 4-65.

How the system selects an interface for incoming packets

The virtual input interface selection for incoming packets is done in the following order:

- 1 If the source IP address located in the packet matches one of virtual interface remote IP addresses (taking into account the netmask), then the matching virtual interface is selected as the input interface.
- 2 If the destination IP address located in the packet matches the virtual interface local IP address (taking into account the netmask), then the matching virtual interface is selected as input interface.
- 3 If there is no match, the input interface will be the main interface.

When virtual interfaces are in use, the statistics are globally counted by the main interface even if the packets are received on virtual interfaces.

Deactivating and deleting WAN virtual IP interfaces

All virtual interfaces are deactivated if the main interface becomes inactive. Deactivating a virtual interface implies a suppression of associated routes.

The interfaces are deleted if one of the following condition occurs:

- The main interface is deleted.
- The lease time has elapsed.

The information about lease time is copied into the virtual interface. Both the shelf and the slot have this information. However, normally the LIM has the responsibility of releasing the interface if the time is elapsed. The copy on the shelf is present if the slot is reset and an update is required.
- A DHCP-RELEASE or DHCP-DECLINE message is relayed by the relay-agent.



Note If you must manually delete a WAN virtual interface on a BIR connection, see “ifmgr” on page A-14.

DHCP relay configuration settings

DHCP relay configuration requires settings in the IP-GLOBAL profile. In addition, LAN and WAN interfaces have settings for enabling or disabling **bootp-relay**. It is enabled by default on LAN interfaces, and disabled by default on WAN interfaces.

Overview of IP-GLOBAL DHCP relay settings

The following IP-GLOBAL parameters (shown with default settings) configure DHCP relay.

```
[in IP-GLOBAL:bootp-relay]
active = no
bootp-servers = [ 0.0.0.0 0.0.0.0 ]
```

Parameter	Setting
active	Enables or disables DHCP relay. When this parameter is set to yes , the IP DSLAM forwards requests from a client on one network (such as a remote interface) to a DHCP server on another network interface.
bootp-servers[1]/[2]	Each of these indexed parameters specifies the IP address of one DHCP server. Only one address is required. When the system receives a DHCP request from a client, the system selects one of these DHCP server address (based on hashing). If a route is not found to the selected DHCP server, the system checks for a route to the second DHCP server address if one is configured. DHCP request packets are not dropped if a route can be found to any of the configured DHCP servers.

Overview of DHCP relay agent settings

All settings within the **relay-agent-information** subprofile (and its subprofiles) apply only when DHCP option 82 is in use. Following are the basic option 82 settings:

```
[in IP-GLOBAL:bootp-relay:relay-agent-information]
enable = no
allow-multiple-interface-creation = no
giaddr-selection = system-ip-address
```

Parameter	Setting
enable	Globally enable or disable DHCP option 82.

Parameter	Setting						
allow-multiple-interface-creation	When set to yes, the relay agent can create a new WAN virtual IP interface on a BIR connection each time a new DHCP request has been successfully serviced. The relay-agent-information field must also be set to yes, and the remote-address field in the BIR CONNECTION profile must be null. (The local-address field can be set to a valid IP address.) If this parameter is set to no, only one interface is associated with a BIR connection.						
giaddr-selection	This parameter specifies which address to use as the gateway address to populate the giaddr field in DHCP packets. The setting applies only when option 82 is configured. When DHCP over BIR is used without option 82, the giaddr field is set to the remote-ip-address configured in the BIR CONNECTION profile. In this case, the DHCP server must know the route to that address. By default, the system address is used as the gateway address, which is also the source address of the DHCP packets. Changing this parameter changes the contents of the giaddr field, not the source address of the packets. Specifying the soft address or local interface address enables the DHCP server to reply to an address rather than the system address. For more details, see “Sample DHCP relay configurations for BIR connections” on page 4-78. Following are valid settings: <table><tr><td>system-ip-address</td><td>Use the configured system IP address (the default). If the system IP address is set to the address of the management interface, the DHCP server must be accessible across that interface.</td></tr><tr><td>soft-ip-address</td><td>Use the IP address configured in the IP-INTERFACE profile with the zero index ({{a a 0 }}), which is the soft interface. If the soft interface is not configured, the system uses the null address (0.0.0.0). The system does not consider virtual interfaces with the zero index ({{a a 0}x}).</td></tr><tr><td>local-ip-address</td><td>Use the IP address of the interface on which the DHCP server is accessible.</td></tr></table>	system-ip-address	Use the configured system IP address (the default). If the system IP address is set to the address of the management interface, the DHCP server must be accessible across that interface.	soft-ip-address	Use the IP address configured in the IP-INTERFACE profile with the zero index ({{a a 0 }}), which is the soft interface. If the soft interface is not configured, the system uses the null address (0.0.0.0). The system does not consider virtual interfaces with the zero index ({{a a 0}x}).	local-ip-address	Use the IP address of the interface on which the DHCP server is accessible.
system-ip-address	Use the configured system IP address (the default). If the system IP address is set to the address of the management interface, the DHCP server must be accessible across that interface.						
soft-ip-address	Use the IP address configured in the IP-INTERFACE profile with the zero index ({{a a 0 }}), which is the soft interface. If the soft interface is not configured, the system uses the null address (0.0.0.0). The system does not consider virtual interfaces with the zero index ({{a a 0}x}).						
local-ip-address	Use the IP address of the interface on which the DHCP server is accessible.						

Overview of DHCP option 82 suboption settings

DHCP option 82 has two suboptions:

- Suboption 1, for conveying a unique identifier for the user circuit
- Suboption 2, for conveying a unique identifier for the remote host

You can configure the system to convey both suboptions or to convey only one of them. Configure suboption 1 in the **circuit-id** subprofile or suboption 2 in the **remote-id** subprofile, or both. If both IDs are enabled, the circuit identifier is followed by the remote identifier.

```
[in IP-GLOBAL:bootp-relay:relay-agent-information:circuit-id]
```

```
enable = no
```

```
send-hostname-only = no
```

```
if-ip = 0.0.0.0
```

```
vendor-option-string = ""
```

```
version = 1
```

```
dhcp-allow-any-src-port = no
```

```
[in IP-GLOBAL:bootp-relay:relay-agent-information:remote-id]
```

```
enable = no
```

```
send-hostname-only = no
```

```
if-ip = 0.0.0.0
```

```
vendor-option-string = ""
```

```
version = 1
```

```
dhcp-allow-any-src-port = no
```

Parameter	Setting

circuit-id remote-id: enable	

	In the circuit-id subprofile, this setting enables or disables the circuit identifier suboption of DHCP option 82. This causes the system to encode the station value (the hostname) of the profile that defines the PVC on which the DHCP client-to-server packet was received and to use it to ensure that DHCP responses are sent to the proper host.

	In the remote-id subprofile, this setting enables or disables the remote identifier suboption of DHCP option 82. This causes the system to encode the remote host end of a circuit on which a DHCP client-to-server packet was received and to use it to ensure that DHCP responses are sent to the proper host.

	If this parameter is set to no , none of the other values within the subprofile are used.

Parameter	Setting
circuit-id remote-id: send-hostname-only	Setting this parameter to yes causes the system to send only the hostname of the subscriber profile (when set in the circuit-id subprofile) or the remote host (when set in the remote-id subprofile) in the DHCP Option 82 suboption field. Other values within the subprofile (such as version or if-ip) are not sent, even if they are configured. Only the actual hostname and its length is sent, not the maximum allowed 32-byte field. <pre> SubOpt Len Sub-option Value +-----+-----+-----+ 1 N HostName +-----+-----+-----+ </pre> For an example, see “Sample configuration that sends only a hostname” on page 4-84.
circuit-id remote-id: if-ip	IP address of one of the control module’s IP interfaces. If both IDs are enabled, only one interface IP address is needed. If this field is empty, the Stinger uses the system address (IP-GLOBAL:system-ip-addr) if that value has been defined. To interoperate with DHCP servers that zero-delimit suboption fields, the specified IP address cannot contain a zero octet.
circuit-id remote-id: vendor-option-string	Currently, this option is supported only for PPPoE snooping. For details, see “Including an identifying text string (PPPoE snooping only)” on page 3-15.
circuit-id remote-id: version	Relay agent version ID. To interoperate with DHCP servers that zero-delimit suboption fields, change this from the default 1 to a value of 257 or higher.
circuit-id remote-id: dhcp-allow-any-src-port	If set to yes , enables the system to process DHCP requests from a nonstandard source port, and to add option 82 to such requests.

DHCP option 82 config sub-options 1 and 2

In earlier versions of the software, the connection name was used as the hostname field in RemoteID suboption. When the connection name was used as the hostname field, both the CircuitID and RemoteID have the same information. A new parameter added in the connection profile, **remote-id**, is now used as the hostname field in the RemoteID of DHCP option 82 (suboption 2). This parameter helps you to provision different hostnames of RemoteID from hostname of CircuitID per connection base.

When **remote-id** is set and **use-extend-remote-id** parameter in the **system** profile is enabled, the **remote-id** value is populated in the hostname field in the remoteID suboption (option 82, suboption 2) for DHCP relay, DHCP snooping, and PPPoE snooping. If **remote-id** is empty, or **use-extend-remote-id** is not enabled, then the value of the name field is used instead. **Remote-id** field allows up to 63 characters.

The `use-extend-remote-id` parameter facilitates backward compatibility. The parameter helps Stinger to switch between the old behavior (use connection name as hostname) and the new behavior (check the `remote-id` field in connection profile). When `use-extend-remote-id` is set to NO, the setting of `remote-id` is ignored and the connection name will be used as the hostname of the `remote-id`. Also, the length of the hostname field in the `remote-id` remains 32 bytes when all fields are sent.

When `use-extend-remote-id` is set to YES, the settings of the `remote-id` decides the hostname field of `remote-id` suboption, and the length of hostname field in `remote-id` holds 64 bytes if all fields are sent. The length of hostname in `circuit-id` (suboption 1) has no change.

According to the RFC 3046, the RemoteID field is unique within the Relay Agent (Stinger). It is not required to check the uniqueness of this field across all connections to avoid cross profile checking. However, the code raises a warning when you try to duplicate the RemoteID.



Note You need to enable the `circuit-id` before you enable `use-extended-remote-id`.

Overview of IP-INTERFACE and CONNECTION DHCP settings

Following are the parameters, shown with default settings, for enabling the use of DHCP across a LAN or WAN interface:

```
[in IP-INTERFACE/{ { shelf-1 first-control-module 1 } 0 }]
bootp-relay-enable = yes

[in CONNECTION/"" :ip-options:bootp-relay-options]
enable = no
max-dynamic-interface = 0
additional-lease-delay = 15

[in CONNECTION/"" :ip-options]
giaddr-assignment = global
```

Parameter	Setting
<code>bootp-relay-enable</code>	Enables or disables relay of DHCP requests and replies received on the LAN interface. The LAN interface on which the DHCP server is accessible must have this parameter set to yes , which is the default setting. Setting this parameter to no prevents relay of DHCP packets received on the interface. It does not prevent the system from sending relayed packets on the interface. This parameter has no impact if the Stinger is either a DHCP client or server.
<code>enable</code>	Enables or disables DHCP relay requests on the interface. This parameter applies to DHCP relay requests on both BIR connections and IP terminated virtual circuits.

Parameter	Setting										
max-dynamic-interface	Maximum number of dynamic interfaces that can be created by the DHCP ACK packet on the connection. If null, dynamic interface creation is disabled on this interface. This parameter applies only to BIR connections.										
additional-lease-delay	Number of seconds to keep the virtual interface up after the lease time (extracted from the DHCP packet, option 51) has expired. A non-null value allows the DHCP client to request a new IP address after the previous IP address lease time has expired without a shutdown of the virtual interface. Typically, the DHCP server reallocates the same IP address to a client. The maximum lease time is 42,949,672 seconds, which translates to 11930 hours (497 days). The Stinger interprets a lease time of 0 as an infinite lease.										
giaddr-assignment	This parameter applies only to BIR connections. Specifies the address to populate the <code>giaddr</code> field in DHCPDISCOVER packets on this connection. Following are valid settings: <table><tr><td>global</td><td>Use the <code>giaddr</code> value specified in the IP-GLOBAL profile. This is the default value.</td></tr><tr><td>system-ip</td><td>Use the configured system IP address.</td></tr><tr><td>soft-ip</td><td>Use the connection's <code>source-if</code> setting. If <code>source-if</code> is null, use the soft interface address.</td></tr><tr><td>local</td><td>Use the local IP address of the interface across which the DHCP server is accessible.</td></tr><tr><td>local-connection</td><td>Use the local IP address of the main BIR connection.</td></tr></table>	global	Use the <code>giaddr</code> value specified in the IP-GLOBAL profile. This is the default value.	system-ip	Use the configured system IP address.	soft-ip	Use the connection's <code>source-if</code> setting. If <code>source-if</code> is null, use the soft interface address.	local	Use the local IP address of the interface across which the DHCP server is accessible.	local-connection	Use the local IP address of the main BIR connection.
global	Use the <code>giaddr</code> value specified in the IP-GLOBAL profile. This is the default value.										
system-ip	Use the configured system IP address.										
soft-ip	Use the connection's <code>source-if</code> setting. If <code>source-if</code> is null, use the soft interface address.										
local	Use the local IP address of the interface across which the DHCP server is accessible.										
local-connection	Use the local IP address of the main BIR connection.										

Per-connection control of the DHCP relay giaddr field

When the Stinger is operating as DHCP relay agent with option 82 configured, the system populates the `giaddr` field in DHCP packets from a client with the address of the relay agent, typically with the Stinger system address or, optionally, another IP address. You can also specify the address used to populate the `giaddr` field on a per-connection basis. If both the global and per-connection settings are used, the per-connection setting takes precedence for a given connection.

Overview of how to use per-connection giaddr contents

When the contents of the `giaddr` field in a DHCP Discover packet is non-zero, DHCP servers choose the pool of addresses from which to allocate the client's address based on the `giaddr` contents. If you configure connections to populate the `giaddr` field with a per-connection address, you can group users into different IP address ranges

and networks. This is a more flexible way to configure addresses via DHCP. The per-connection assignment of a relay-agent address also facilitates growth. When you add a new LIM to a Stinger system, a new address range can be designated for users of that LIM without any reconfiguration of the DHCP server.

Samples of how various settings affect giaddr contents

Table 4-3 shows how sample CONNECTION profile and IP-INTERFACE profile settings affect giaddr contents. For a sample configuration that controls the giaddr field contents on a per-connection basis, see “Sample using the soft interface address in the giaddr field” on page 4-82.

Table 4-3. Settings that affect contents of the giaddr field in DHCP packets

CONNECTION settings	Address used in giaddr field	Description
set ip-options giaddr-assignment = soft-ip	new ip-interface { { 0 0 0 } 2 }	Populating the giaddr field with a soft interface address.
set ip-options source-if = sip2	set ip-address = 1.1.1.254	
set ip-options giaddr-assignment = local	new ip-interface { { 1 8 2 } 0 } set ip-address = 10.10.10.10	Populating the giaddr field with the IP address of the local interface (typically a GigE interface) across which packets to the DHCP server are routed.
set ip-options giaddr-assignment = local-connection	set ip-options local-address = 10.10.1.2	Populating the giaddr field with the local IP address of the CONNECTION profile.


Sample DHCP relay configurations for IPoA connections

To support centralized assignment of the IP addresses of host PCs via DHCP, both the remote CPE router and the IP DSLAM must be configured as DHCP relay agents. The sample configurations in this section show how to use DHCP relay with or without option 82.

Sample configuration using DHCP relay without option 82

Figure 4-18 shows a high-speed circuit to a remote CPE router with a LAN interface to one or more host PCs.

Figure 4-18. DHCP relay for an IPoA terminated PVC



In this example, the remote PC is configured to use DHCP to obtain an IP address. The CPE router on the remote LAN is configured statically with the IP address 3.3.3.3 and is configured to enable DHCP relay. The CPE router DHCP relay configuration

specifies a DHCP server at 2.2.2.142 (the DHCP server beyond the DHCP relay agent on the IP DSLAM).

The IP DSLAM is configured as DHCP relay agent with the DHCP server at 2.2.2.142 across its Gigabit Ethernet interface (2.2.2.2). The DHCP server is configured to recognize the CPE router at 3.3.3.3.

The following commands configure the IP DSLAM as a DHCP relay agent:

```
admin> read ip-global
admin> set bootp-relay active = yes
admin> set bootp-relay bootp-servers 1 = 2.2.2.142
admin> write -f
```

The next commands configure the CONNECTION profile for the CPE router:

```
admin> new connection cpe-router
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 3.3.3.3/30
admin> set ip-options bootp-relay-options enable = yes
admin> set atm-options nailed-group = 251
admin> write -f
```

Sample configuration using DHCP relay with option 82

This example builds on the sample DHCP relay configuration shown above (Figure 4-18). The CONNECTION profile to the CPE router in that sample configuration does not require any changes to support option 82.

For sites that support option 82, the DHCP server configuration typically requires the presence of an ID in DHCP queries. For example, the DHCP server in this example is configured to recognize the CPE router at 3.3.3.3 across interface 7.7.7.7, and to require a circuit ID. If the DHCP request forwarded to the server by the IP DSLAM does not contain the circuit ID, the server refuses to return an address.

The following commands configure the system to enable the circuit identifier suboption of DHCP option 82 and specify the Gigabit Ethernet address as the ID (Figure 4-18):

```
admin> read ip-global
admin> set bootp-relay relay-agent-information circuit-id enable = yes
admin> set bootp-relay relay-agent-information circuit-id if-id = 2.2.2.2
admin> write -f
```

Sending only a hostname in the suboption fields


The `send-only-hostname` identifier option can be used for both IPoA and BIR connections. For an example configuration that shows a BIR connection, see “Sample configuration that sends only a hostname” on page 4-84.

Interoperation with DHCP servers that zero-delimit suboption fields

Some DHCP servers treat option-82 suboption information fields as zero-delimited strings rather than as an array of bytes. When a suboption information field (`circuit-id` and `remote-id`) contains a zero byte in the middle of the field, these DHCP servers consider the zero byte to be the end of the option-82 string.

Figure 4-19 shows the format Stinger systems use for the DHCP option-82 suboption information field (`circuit-id` and `remote-id`):

Figure 4-19. Option 82 information field formats



The two-byte version number of the relay-agent information ID is set by default to 1, which is coded internally as 0x00 0x01.

The four-byte relay-agent IP address is set to an IP address in use in the system.

To avoid conflicts with DHCP servers that zero-delimit suboption fields, the version number setting is configurable by the operator. In addition, operators must be sure not to assign an IP address that contains a zero byte. For example, the following commands set the circuit identifier suboption version field to 257:

```
admin> read ip-global
admin> set bootp-relay relay-agent-information circuit-id version = 257
admin> write -f
```

Allowing non-standard DHCP source ports

To enable the Stinger to process DHCP requests from a nonstandard source port, and to add option 82 (which includes the source port number) to such requests, you can configure the system to allow any DHCP source port. For example:

```
admin> read ip-global
admin> set bootp-relay relay-agent circuit-id dhcp-allow-any-src-port = yes
admin> write -f
```

DHCP issues on LAN management interfaces

The management interface of the Stinger unit is the IP interface used to access the system for management or configuration tasks, usually via `telnet` or an SNMP manager. It might be an 10/100 Ethernet port, a Gigabit Ethernet port, or an ATM trunk interface.


Often the system address is set to the IP address of the management interface. By default, this is also the address that populates the `giaddr` field in DHCP packets the system relays out on a LAN interface.

If the system address is set to the management interface and the DHCP server is not accessible across that interface, you should specify either the soft IP address or the local interface address to populate the `giaddr` field of DHCP packets. This enables the server to reply to an address rather than the system address.

You can also disable forwarding of DHCP packets received on the management network interface, or any LAN interface.

The sample configuration in this section uses the network setup and IP addresses shown in Figure 4-20:

Figure 4-20. Sample DHCP usage with LAN management interface



The following commands configure the management network interface. DHCP relay is blocked on this interface.

```
admin> read ip-interface { { 1 first-control-module 1 } 0 }
admin> set ip-address = 10.66.66.30/8
admin> set bootp-relay-enable = no
admin> write -f
```

The following commands configure the Gigabit Ethernet network interface. This is the interface on which the DHCP server is accessible, so DHCP relay must be enabled (as it is by default).

```
admin> read ip-interface { { 1 first-control-module 2 } 0 }
admin> set ip-address = 192.168.144.200/24
admin> write -f
```

The following commands set the system address to the management network address, configure the DHCP server address, and set the `giaddr` address to the local interface on which the DHCP server is accessible.

```
admin> read ip-global
admin> set system-ip-addr = 10.66.66.30/8
admin> set bootp-relay active = yes
admin> set bootp-relay bootp-servers 1 = 192.168.144.1
```

```
admin> set bootp-relay relay-agent enable = yes
admin> set bootp-relay relay-agent giaddr-selection = local-ip-address
admin> write -f
```

Sample DHCP relay configurations for BIR connections


The sample configurations in this section are intended to show how the system behaves with specific global and connection settings.



Note In each sample configuration, connection-specific configurations other than the settings shown are meaningless and forbidden.

All configurations in this section use the basic network setup and IP addresses shown in Figure 4-21:

Figure 4-21. Sample DHCP setup on BIR connection



Sample configuration with no DHCP relay on a BIR connection

In this sample configuration, the following conditions apply:

- Relay-agent is disabled on this interface. If the system receives DHCP requests on the interface, it discards them.
- Each remote host on the BIR connection has a static IP configuration.

The following set of commands configures the system to act as DHCP relay agent, but does not enable option 82 or multiple interface creation on BIR connections:

```
admin> read ip-global
admin> set bootp-relay active = yes
admin> set bootp-relay bootp-servers 1 = 192.168.144.1
admin> set bootp-relay relay-agent-information enable = no
admin> set bootp-relay relay-agent-information allow-multiple-interface = no
admin> set bootp-relay relay-agent-information circuit-id enable = no
admin> set bootp-relay relay-agent-information circuit-id if-ip = 0
admin> write -f
```

The following commands configure the BIR connection and disable DHCP relay requests:

```
admin> read connection
admin> set station = bir-host-1
admin> set active = yes
```



```
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 192.168.146.10/24
admin> set ip-options local-address = 192.168.146.200/24
admin> set ip-options bootp-relay-options enable = no
admin> set ip-options bootp-relay-options max-dynamic-interface = 0
admin> set atm-options nailed-group = 3
admin> set bir-options enable = yes
admin> write -f
```

Sample configuration enabling relay agent on a BIR connection

In this sample configuration, the following conditions apply:

- Relay-agent is enabled on this interface. DHCP option 82 is not enabled.
- The `giaddr` field of the DHCP packet will contain the `local-address` of the interface on which the request was made. This address will be used to route back the DHCP reply.
- Several hosts acting as DHCP clients can be configured on this BIR connection, but each host must be on the same IP network (for example, 192.168.146.0).
- All hosts share the same gateway located at 192.168.146.200.
- Some hosts on this connection can have a static IP configuration.
- The Stinger has only one interface associated with the BIR connection.

The following set of commands configures the system to act as DHCP relay agent, but does not enable option 82 or multiple interface creation on BIR connections.

```
admin> read ip-global
admin> set bootp-relay active = yes
admin> set bootp-relay bootp-servers 1 = 192.168.144.1
admin> set bootp-relay relay-agent-information enable = no
admin> set bootp-relay relay-agent-information allow-multiple-interface = no
admin> set bootp-relay relay-agent-information circuit-id enable = no
admin> set bootp-relay relay-agent-information circuit-id if-ip = 0
admin> write -f
```

The following commands enable relay agent requests on the BIR connection, but do not enable option 82.

```
admin> read connection
admin> set station = bir
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 192.168.146.10/24
admin> set ip-options local-address = 192.168.146.200/24
admin> set ip-options bootp-relay-options enable = yes
admin> set ip-options bootp-relay-options max-dynamic-interface = 0
admin> set atm-options nailed-group = 3
```

```
admin> set bir-options enable = yes
admin> write -f
```

Sample configuration with option 82

In this sample configuration, the following conditions apply:

- Relay-agent is enabled on this interface.
- DHCP option 82 is enabled. The `giaddr` field of the DHCP packet will contain the Stinger `system-ip-address`, and the option 82 identifier will be inserted in the DHCP packet. Option 82 contains enough information to route back the DHCP reply.
- Several hosts acting as DHCP clients can be configured on this BIR connection.
- All hosts share the same gateway located at 192.168.146.200.
- Some hosts on this connection can have a static IP configuration.
- The Stinger has only one interface associated with the BIR connection.

The following set of commands configures the system to act as DHCP relay agent, and enables DHCP option 82 in the system. The option 82 identifier inserted into DHCP packets will be associated with the virtual circuit to the remote device.

```
admin> read ip-global
admin> set bootp-relay active = yes
admin> set bootp-relay bootp-servers 1 = 192.168.144.1
admin> set bootp-relay relay-agent-info enable = yes
admin> set bootp-relay relay-agent-info allow-multiple-interface = no
admin> set bootp-relay relay-agent-info circuit-id enable = yes
admin> set bootp-relay relay-agent-info circuit-id if-ip = 192.168.146.200
admin> write -f
```

The following commands enable relay agent requests on the BIR connection:

```
admin> read connection
admin> set station = bir
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 192.168.146.10/24
admin> set ip-options local-address = 192.168.146.200/24
admin> set ip-options bootp-relay-options enable = yes
admin> set ip-options bootp-relay-options max-dynamic-interface = 0
admin> set atm-options nailed-group = 3
admin> set bir-options enable = yes
admin> write -f
```

Sample configuration with option 82 and multiple interface creation

In this sample configuration, the following conditions apply:

- Relay-agent is enabled on this interface.
- DHCP option 82 is enabled. The `giaddr` field of the DHCP packet will contain the Stinger `system-ip-address`, and the option 82 ID will be inserted in the DHCP packet. Option 82 contains enough information to route back the DHCP reply.
- Several hosts acting as DHCP clients can be configured on this BIR connection.
- All hosts share the same gateway located at 192.168.146.200.
- Static IP configuration for hosts on the BIR connection is forbidden.
- The Stinger has the same number of interfaces associated with the BIR connection as there are remote hosts. In this example, the BIR connection supports up to three WAN virtual interfaces for three DHCP clients.

The system creates a new interface for each fully serviced DHCP request. In this sample configuration, the IP address is obtained from the `local-address` defined in the CONNECTION profile of the configured main interface, and the remote address is obtained dynamically as described in “IP address assignments on DHCP-configured WAN virtual interfaces” on page 4-66. Other Interface parameters are cloned from the configured main interface. The system sets up direct routes to the virtual interfaces during interface creation, and assigns the type `virtual` to the new interfaces.

The following set of commands configures the system to act as DHCP relay agent, enables option 82, and enables creation of virtual interfaces for BIR connections.

```
admin> read ip-global
admin> set bootp-relay active = yes
admin> set bootp-relay bootp-servers 1 = 192.168.144.1
admin> set bootp-relay relay-agent-info enable = yes
admin> set bootp-relay relay-agent-info allow-multiple-interface = yes
admin> set bootp-relay relay-agent-info circuit-id enable = yes
admin> set bootp-relay relay-agent-info circuit-id if-ip = 192.168.146.200
admin> write -f
```

The next commands enable relay agent requests on the BIR connection. Note that the `remote-address` field must be null. For details about addresses assigned to the WAN virtual IP interfaces as they are created, see “IP address assignments on DHCP-configured WAN virtual interfaces” on page 4-66.

```
admin> read connection
admin> set station = bir
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 0.0.0.0/0
admin> set ip-options local-address = 192.168.146.200/24
admin> set ip-options bootp-relay-options enable = yes
admin> set ip-options bootp-relay-options max-dynamic-interface = 3
```

```
admin> set atm-options nailed-group = 3
admin> set bir-options enable = yes
admin> set bir-options proxy-arp = yes
admin> write -f
```

Sample configuration on a VDSL port with option 82 and multiple interface creation

On very high speed DSL (VDSL) ports, BIR connections use Ethernet encapsulation on the drop. In this example, DHCP option 82 is enabled, and several hosts acting as DHCP clients can be configured on this BIR connection.

The following set of commands configures the system to act as DHCP relay agent, enables option 82, and enables creation of virtual interfaces for BIR connections.

```
admin> read ip-global
admin> set bootp-relay active = yes
admin> set bootp-relay bootp-servers 1 = 192.168.144.1
admin> set bootp-relay relay-agent-info enable = yes
admin> set bootp-relay relay-agent-info allow-multiple-interface = yes
admin> write -f
```

The following commands enable relay agent requests for the BIR connection on a VDSL port:

```
admin> new connection
admin> set station = bir-vdsl-dhcp
admin> set active = yes
admin> set encapsulation-protocol = ethernet
admin> set ip-options remote-address = 0.0.0.0
admin> set ip-options bootp-relay-options enable = yes
admin> set ip-options bootp-relay-options max-dynamic-interface = 3
admin> set bir-options enable = yes
admin> set bir-options proxy-arp = yes
admin> set ether-options nailed-group = 251
admin> write -f
```

For background information about the ETHER-QOS profile and Ethernet options in CONNECTION profiles, see the *Stinger VDSL Line Interface Module (LIM) Guide*.

Sample using the soft interface address in the giaddr field

In this sample configuration, a second soft interface is created and its address is used to populate the giaddr field in DHCP Discover packets. The DHCP server is configured so that the IP address assigned to the client is in the same subnet as the IP address in the giaddr field.

- 1 Enable DHCP Relay with option 82. For example:

```
admin> read ip-global
admin> set bootp-relay active = yes
```

```
admin> set bootp-relay bootp-servers 1 = 11.11.11.45
admin> set bootp-relay relay-agent-information enable = yes
admin> set bootp-relay relay-agent allow-multiple-interface = yes
admin> set bootp-relay relay-agent circuit-id enable = yes
admin> set bootp-relay relay-agent circuit-id if-ip = 2.2.2.222
admin> set bootp-relay relay-agent circuit-id dhcp-allow-any = yes
admin> set bootp-relay relay-agent giaddr-selection = local-ip-address
admin> write -f
```

2 Create a soft interface and assign an IP address. For example:

```
admin> new ip-interface
admin> set ip-address = 22.22.22.5
admin> write -f
```

3 Create a CONNECTION profile for a BIR client, and specify that the soft IP interface address should be used to populate the giaddr field. For example:

```
admin> new connection bir-client
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options bootp-relay-options enable = yes
admin> set ip-options bootp-relay-options max-dynamic-interface = 7
admin> set ip-options giaddr-assignment = soft-ip
admin> set ip-options local-address = 33.33.33.31/8
admin> set atm-options vci = 66
admin> set atm-options nailed-group = 67
admin> set bir-options enable = yes
admin> set bir-options proxy-arp = yes
admin> write -f
```

Sample configuration using the DHCP router option

This sample configuration is just like the previous one (“Sample configuration with option 82 and multiple interface creation” on page 4-81) except that the IP address (the local address) of the virtual interface is obtained from the router option (option 3). Using the DHCP router option enables a host to use a particular gateway, and so allows the DHCP server to distribute addresses on different subnets for hosts on the same physical connection.

If the information is not provided in the DHCP packet, the system discards the DHCP reply and does not create the WAN virtual interface. If more than one address is specified via the router option in the reply packet, the system uses the first one.

The following set of commands configures the system to act as DHCP relay agent, enables option 82, and enables creation of virtual interfaces for BIR connections.

```
admin> read ip-global
admin> set bootp-relay active = yes
admin> set bootp-relay bootp-servers 1 = 192.168.144.1
```

```
admin> set bootp-relay relay-agent-info enable = yes
admin> set bootp-relay relay-agent-info allow-multiple-interface = yes
admin> set bootp-relay relay-agent-info circuit-id enable = yes
admin> set bootp-relay relay-agent-info circuit-id if-ip = 192.168.146.200
admin> write -f
```

The next commands enable relay agent requests on the BIR connection. Note that the `remote-address` field must be null. For details about addresses assigned to the virtual IP interfaces as they are created, see “IP address assignments on DHCP-configured WAN virtual interfaces” on page 4-66.

```
admin> read connection
admin> set station = bir
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 0.0.0.0/0
admin> set ip-options local-address = 0.0.0.0/0
admin> set ip-options bootp-relay-options enable = yes
admin> set ip-options bootp-relay-options max-dynamic-interface = 3
admin> set atm-options nailed-group = 3
admin> set bir-options enable = yes
admin> write -f
```

Sample configuration that sends only a hostname

In this sample configuration, DHCP relay with option 82 is enabled and several remote hosts acting as DHCP clients are configured. For the purposes of this example, only the `circuit-id` suboption is used. However, with this setup either the `circuit-id` or `remote-id` suboption could be used, singly or together.

The following commands configure DHCP relay options in the IP-GLOBAL profile. With these settings, only the hostname will be sent in the `circuit-id` suboption field:

```
admin> read ip-global
admin> set bootp-relay active = yes
admin> set bootp-relay bootp-servers 1 = 192.168.144.1
admin> set bootp-relay relay-agent-info enable = yes
admin> set bootp-relay relay-agent-info circuit-id enable = yes
admin> set bootp-relay relay-agent-info circuit-id send-hostname-only = yes
admin> set bootp-relay relay-agent-info giaddr-selection = local-ip-address
admin> write -f
```

The following commands configure a local routed VLAN interface:

```
admin> new ip-interface { { 1 8 2 } 500 }
admin> set ip-address = 192.168.146.200/24
admin> set vlan-enabled = yes
admin> set vlan-id = 500
```

```
admin> write -f
admin>
admin> new vlan-ethernet { { 1 8 2 } 500 }
admin> set enabled = yes
admin> write -f
```

The following commands configure the BIR connection to obtain its configuration via DHCP:

```
admin> new connection
admin> set station = bridge-cpe
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options bootp-relay-options enable = yes
admin> set ip-options bootp-relay-options max-dynamic-interface = 7
admin> set telco-options nailed-groups = 711
admin> set atm-options vpi = 2
admin> set atm-options vci = 36
admin> set atm-options nailed-group = 711
admin> set bir-options enable = yes
admin> write -f
```

Configuring broadband RAS subscriber access

When it is acting as a broadband remote access server (BRAS), a Stinger IP DSLAM can terminate PPPoA and PPPoE sessions received on a LIM DSL interface.

For information about bridging traffic to an upstream BRAS for termination, see Chapter 3, “VLAN Configuration.”

When a DSL interface receives a PPPoA or PPPoE packet stream, the system switches the inbound traffic to the LIM’s ATM internal interface on the basis of an ATM circuit configuration. The IP control module then uses IP routing to forward the packet streams to the appropriate egress interface. For background information about ATM circuit configurations, see the *Stinger ATM Configuration Guide*.



Note PPPoA and PPPoE sessions are supported on DSL interfaces only. PPPoA and PPPoE are not supported on T1/E1/IMA ports or across trunk or Ethernet interfaces.

Recommended call-type setting for PPP sessions

The system assigns nailed group 1 to LIM slot 1 port 1 ({ 1 1 1 }), and also assigns nailed group 1 to permanent, leased PPP connections. Because the default call-type in Stinger systems is nailed (ft1), a conflict occurs when you configure PPPoA or PPPoE connections on any LIM interface, and { 1 1 1 } is enabled.

Following are representative default settings showing the conflicting nailed group assignments:

```
[in DS1-ATM/{ shelf-1 slot-1 1 }:line-config]
```

```
nailed-group = 1
[in CONNECTION/"" :telco-options]
call-type = ft1
nailed-groups = 1
```



Note Because these defaults conflict, it is recommended to set the **call-type** value to off for all PPP connections.


With the default **call-type** setting of **ft1** in a PPP profile, if { 1 1 1 } is enabled (with or without connections), the system generates rolling LOG errors, LOG warning, and LOG information messages. This condition is prevented by setting **call-type** to off.

In addition, when **call-type** is set to off, the system is able to terminate PPP sessions automatically following a period of client inactivity, and reestablish them when the client becomes active again.

Overview of PPPoA and PPPoE topologies

A PPPoA connection uses the ATM adaptation layer 5 (AAL5) protocol as a framing mechanism across point-to-point virtual circuits, as described in RFC 2364, *PPP over AAL5*. Only VC-multiplexed PPPoA is currently supported. The PPPoA session is between a router and a Stinger unit, as shown in Figure 4-22. The PPPoA session enables an IP client on the far side of the PPPoA router to connect through the Stinger unit to the IP cloud beyond it.


Figure 4-22. PPPoA topology



For PPPoE, the connection uses Ethernet-bridged framing, as defined in RFC 2516, *A Method for Transmitting PPP Over Ethernet (PPPoE)*. PPPoE is negotiated in two phases: a discovery phase and a session phase. During the discovery phase, the PPPoE client negotiates with the Stinger to obtain information it requires. When the session has been established, the client sends PPP packets encapsulated in Ethernet-bridged frames.

The PPPoE sessions shown in Figure 4-23 are between a PPPoE client connected through the bridge, and the Stinger unit. The PPPoE session enables the client to connect through the Stinger unit to the IP cloud beyond it.

Figure 4-23. PPPoE topology



Required setup for PPPoA and PPPoE connections

To enable establishment of PPPoA or PPPoE connections using IP routing, you must complete the following steps:

1. Configure the ANSWER-DEFAULTS profile to accept PPP session requests and require their authentication.
You do this once, to enable the system to process all subsequent PPP session requests. See “Configuring the ANSWER-DEFAULTS profile for PPP sessions” on page 4-87.
2. Configure a CONNECTION profile for the underlying ATM circuit required between the DSL interface on which the PPP connection will be established and the LIM internal interface.
 - For PPPoE, the underlying circuit must use RFC 1483 LLC.See “Terminating traffic on a LIM internal interface” on page 4-89.
3. Configure a CONNECTION or RADIUS profile for each PPPoE or PPPoA session.
See “Example of configuring a PPPoA connection” on page 4-91 or “Example of configuring a PPPoE connection” on page 4-95.

If you are using RADIUS to externally authenticate PPP sessions, see the *TAOS RADIUS Guide and Reference*.

Configuring the ANSWER-DEFAULTS profile for PPP sessions

To control access, PPP sessions typically require password authentication each time a PPP session is opened. Stinger units support RFC 1334 *PPP Authentication Protocol (PAP)* and RFC 1994 *Challenge Handshake Authentication Protocol (CHAP)* authentication, which must be negotiated with the client on the basis of settings in the ANSWER-DEFAULTS profile and the client’s CONNECTION or RADIUS profile.

The values in the ANSWER-DEFAULTS profile are applied before the system locates the CONNECTION or RADIUS profile associated with the request. If the PPP client’s profile contains a similar parameter with a different value, the connection-specific value overrides the ANSWER-DEFAULTS value when building the session. If no similar value is found in the CONNECTION or RADIUS profile, the ANSWER-DEFAULTS setting is used.

Following are some ANSWER-DEFAULTS parameters, shown with default settings, that affect PPP authentication and session timeout:

```
[in ANSWER-DEFAULTS]
profiles-required = yes
[in ANSWER-DEFAULTS:ppp-answer]
```

```

enabled = yes
receive-auth-mode = no-ppp-auth
bi-directional-auth = none
substitute-send-name = ""
lcp-keepalive-period = 0
lcp-missed-keepalives = 4
use-magic-number = no

[in ANSWER-DEFAULTS:session-info]
idle-timer = 120
max-call-duration = 0

```

Parameter	Setting
profiles-required	A setting of yes (the default) prevents unauthenticated sessions. If set to no , the system builds a temporary profile for session requests for which it cannot locate a configured profile.
enabled	The enabled parameter must be set to yes (the default) for the system to answer PPP session requests.
receive-auth-mode	With the default no-ppp-auth setting, the Stinger unit does not request authentication. If set to a non-default value, the Stinger unit requests an authentication protocol, and the client must accept one of the options the system offers.
bi-directional-auth	Support for bidirectional CHAP. If set to allowed or required , the system negotiates bidirectional CHAP if the client's CONNECTION profile specifies the proper settings.
substitute-send-name	System name to send to clients for bidirectional CHAP authentication, if different from the name setting in the SYSTEM profile.
lcp-keepalive-period	See "Enabling LCP keepalives" on page 4-98.
lcp-missed-keepalives	See "Enabling LCP keepalives" on page 4-98.
use-magic-number	See "Enabling magic-number negotiation when LQM is disabled" on page 4-101.
idle-timer	With a call-type setting of off in a client profile, the system uses the idle-timer value to terminate the session after a default interval of 2 minutes. You can configure a different default interval here or in the client's session-options subprofile, by specifying the maximum number of consecutive seconds a session can remain idle before it is terminated.
max-call-duration	Maximum number of minutes of connect time for a PPP session. The default zero value disables the timer.

Most sites change the default setting of the **receive-auth-mode** parameter to ensure authentication of a PPP request before a session can be established. For example:

```

admin> read answer-defaults
admin> set ppp-answer receive-auth-mode = any-ppp-auth

```

```
admin> write -f
```

With this setting, the system accepts session requests that provide any of the supported PPP authentication methods, but it drops requests that do not offer any authentication protocols during session negotiation.

The following commands enable bidirectional authentication for sessions that use CHAP and specify the proper settings in the CONNECTION or RADIUS profile:

```
admin> set ppp-answer bidirectional-auth = allowed
admin> write -f
```

With these settings, if a calling device accepts CHAP authentication, the system attempts to negotiate bidirectional CHAP, but does not reject the request if the negotiation fails. However, if bidirectional CHAP is negotiated, authentication must succeed in both directions. For related information, see “Sample PPPoA connection with bidirectional CHAP authentication” on page 4-93.

Terminating traffic on a LIM internal interface

For all installed LIMs that can terminate PPPoA or PPPoE calls, the system creates an ATM-INTERNAL profile for the LIM’s internal ATM segmentation assembly and reassembly (SAR) port. The internal interface number is one greater than the highest DSL interface number on the module. For example, the following command output indicates a 48-port LIM in slot 1, a 72-port LIM in slots 2 and 4, and a 24-port LIM in slot 6:

```
admin> dir atm-internal
42 05/09/2005 08:41:34 { shelf-1 slot-1 49 } 1:1:49
42 05/09/2005 08:41:34 { shelf-1 slot-2 73 } 1:2:73
42 05/09/2005 09:30:38 { shelf-1 slot-4 73 } 1:4:73
42 05/09/2005 08:41:34 { shelf-1 slot-6 25 } 1:6:25
38 05/09/2005 08:41:30 { shelf-1 first-control-module 1 } 1:8:1
```

The underlying ATM circuit required to terminate traffic

For a DSL interface to handle PPPoA and PPPoE incoming calls, you must configure an ATM circuit connection between the external DSL interface and the LIM’s internal ATM interface. The ATM circuit configuration must specify the nailed group of the internal interface as the second leg of the circuit (in the `atm-connect-options` subprofile).

For background information about configuring ATM circuits, see the *Stinger ATM Configuration Guide*. For an overview of the relevant settings, see “Typical CONNECTION atm-options settings for terminating PVCs” on page 4-49.



Note The system prevents configuration of a circuit from one LIM to the ATM internal interface of another LIM, or configurations that attempt to use the internal interface of a LIM in other ways.

PPPoA and PPPoE encapsulation types in the underlying ATM circuit

The `atm1483type` setting in the CONNECTION profile for the ATM circuit connection must be set properly for the PPPoA or PPPoE connection to come up. Table 4-4 shows the required settings:

Table 4-4. Required encapsulation types for PPPoA and PPPoE

Connection type	CONNECTION setting	Description
PPPoA	set atm-options atm1483type = aal5-vc	The system can establish a PPPoA connection regardless of the atm1483type setting in the CONNECTION profile, but the CPE must be configured to do VC multiplexing. For this reason, the aal5-vc setting in the CONNECTION profile is recommended for PPPoA connections.
PPPoE	set atm-options atm1483type = aal5-llc	The system requires LLC encapsulation both in the atm1483type setting in the CONNECTION profile and in the CPE configuration. The aal5-llc setting in the CONNECTION profile is required for PPPoE connections. This is the default setting.
PPP autodetection	set atm-options atm1483type = aal5-llc	For PPP autodetection to work properly, the aal5-llc setting of the atm1483type parameter in the CONNECTION profile is required. This is the default setting.

Configuring an ATM circuit to terminate PPPoA

The following commands enable an SDSL interface in slot 4 and configure an ATM circuit from the DSL interface to the LIM's ATM internal interface.

To determine the nailed group of a LIM's internal ATM interface, use the `which` command. Following is a sample command for the internal interface of a 72-port LIM in slot 4:

```
admin> which -n { 1 4 73 }
```

Nailed group corresponding to port { shelf-1 slot-4 73 } is 2271

The following ATM circuit uses aal5-vc encapsulation, which is required for PPPoA termination.

```
admin> read sdsl { 1 4 1 }
```

```
admin> set enabled = yes
```

```
admin> write -f
```

```
admin> new connection cir-1
```

```
admin> set active = yes
```

```
admin> set encapsulation-protocol = atm-circuit
```

```
admin> set atm-options vci = 38
```

```
admin> set atm-options nailed-group = 151
```

```
admin> set atm-options atm1483type = aal5-vc
```

```
admin> set atm-connect-options vci = 36
```

```
admin> set atm-connect-options nailed-group = 2271
```

```
admin> set atm-connect-options atm1483type = aal5-vc
```

```
admin> write -f
```

Configuring an ATM circuit to terminate PPPoE

The following commands enable an SDSL interface in slot 7 and configure an ATM circuit from the DSL interface to the LIM's ATM internal interface.

To determine the nailed group of a LIM's internal ATM interface, use the `which` command. Following is a sample command for the internal interface of a 72-port LIM in slot 5:

```
admin> which -n { 1 5 73 }  
Nailed group corresponding to port { shelf-1 slot-4 73 } is 2321
```

The following ATM circuit uses `aal5-llc` encapsulation, which is strongly recommended for PPPoE termination.


```
admin> read sdsl { 1 5 1 }  
admin> set enabled = yes  
admin> write -f  
admin> new connection cir-2  
admin> set active = yes  
admin> set encapsulation-protocol = atm-circuit  
admin> set atm-options vci = 45  
admin> set atm-options nailed-group = 201  
admin> set atm-options atm1483type = aal5-llc  
admin> set atm-connect-options vci = 47  
admin> set atm-connect-options nailed-group = 2321  
admin> set atm-connect-options atm1483type = aal5-llc  
admin> write -f
```

For information about applying a PPPoE server configuration to the ATM circuit profile, see “Enabling support for PPPoE service names” on page 4-101.

Example of configuring a PPPoA connection

Figure 4-24 shows a Stinger IP DSLAM. Across a DSL interface, a CellPipe® unit is operating as a PPPoA router.

Figure 4-24. Example of a PPPoA session on a DSL interface



Note You can use bidirectional CHAP authentication for PPPoA connection, but it is not required. The sample configuration provides an example of how to use bidirectional authentication.

Overview of PPPoA CONNECTION settings

For background information about IP routing configurations, see “Configuring IPoA subscriber connections” on page 4-49. In addition to those settings, following are relevant PPPoA parameters, shown with default settings, including bidirectional CHAP authentication:

```
[in CONNECTION/""]  
station* = ""  
encapsulation-protocol = atm-circuit  
  
[in CONNECTION/"":ip-options]  
ip-routing-enabled = yes  
remote-address = 0.0.0.0/0  
  
[in CONNECTION/"":ppp-options]  
send-auth-mode = none  
bidirectional-auth = none  
substitute-recv-name = ""  
send-password = ""  
recv-password = ""  
  
[in CONNECTION/"":pppoa-options]  
pppoa = yes  
  
[in CONNECTION/"":telco-options]  
call-type = ft1
```

Parameter	RADIUS attribute	Setting
station	User-Name (1)	Name of the PPPoA router. The value is case sensitive and must exactly match the name the router sends during authentication.
encapsulation-protocol	Framed-Protocol (7)	Encapsulation protocol. Set to ppp for PPPoA clients.
ip-routing-enabled	Ascend-Route-IP (228)	Enables or disables IP routing for the interface. IP routing is enabled by default.
remote-address	Framed-IP-Address (8) Framed-IP-Netmask (9)	IP address of the PPPoA router, which can include a subnet specification.
bidirectional-auth	Ascend-Bi-Directional-Auth (46)	Enables or disables bidirectional CHAP authentication. Used only for bidirectional CHAP.
send-auth-mode	Ascend-Send-Auth (231)	Set to chap-ppp-auth for bidirectional CHAP. Used only for bidirectional CHAP.
substitute-recv-name	Ascend-Recv-Name (45)	Name that must be received from the far end during bidirectional CHAP authentication, if different from the station setting. Used only for bidirectional CHAP.
send-password	Ascend-Send-Secret (214)	Password the Stinger must send to the far end during bidirectional CHAP, used only for bidirectional CHAP.
recv-password	Password (2)	Password the system must receive from the PPPoA router. This setting is used for all PPP authentication methods.

Parameter	RADIUS attribute	Setting
pppoa	Ascend-PPPoA-Enable (318)	Enables or disables processing of PPPoA packets. PPPoA is enabled by default. If both pppoe and pppoa are set to yes in the same CONNECTION profile, the system automatically detects which protocol is used in the incoming traffic stream and performs the required processing.
call-type	Ascend-Call-Type (177)	Set call-type to off in PPP profiles. For background information, see “Recommended call-type setting for PPP sessions” on page 4-85.

Sample PPPoA connection with bidirectional CHAP authentication

First, configure an ATM circuit between the DSL interface connecting to the PPPoA router and the LIM's internal interface. See “Terminating traffic on a LIM internal interface” on page 4-89.

Then, verify that the ANSWER-DEFAULTS profile enables the PPP authentication methods to be used, and configure a CONNECTION profile to the PPPoA router.

For example, the following commands configure a profile in which bidirectional CHAP authentication is required with the CellPipe® router in Figure 4-24:

```
admin> new connection cellpipe1
admin> set active = yes
admin> set encapsulation-protocol = ppp
admin> set ip-options remote-address = 3.3.3.141/29
admin> set ppp-options send-auth-mode = chap-ppp-auth
admin> set ppp-options bi-directional-auth = required
admin> set ppp-options send-password = sendpw
admin> set ppp-options recv-password = recvpw
admin> set telco-options call-type = off
admin> write -f
```

Following is a comparable RADIUS user profile:

```
cellpipe1 Password = "recvpw"
  Service-Type = Framed-User,
  Ascend-Require-Auth = Require-Auth,
  Ascend-Auth-Type = Auth-CHAP,
  Ascend-Send-Auth = Send-Auth-CHAP,
  Ascend-Bi-Directional-Auth=Bi-Directional-Auth-Required,
  Ascend-Send-Secret = "sendpw",
  Ascend-Call-Type = Switched,
  Framed-Protocol = PPP,
  Framed-IP-Address = 3.3.3.141,
  Framed-IP-Netmask = 255.255.255.248,
```

PPPoA over LLC

To bring up a PPPoA call, Stinger requires two profiles, atm-ckt and PPPoA. The atm-ckt profile runs between the dsl port and the sar port. Both VC-mux and LLC type of encapsulations use this approach.

The first packet that the circuit receives decides the type of encapsulation to be used. For example,

- If the first packet arrives with a VC encapsulation, then Stinger brings up the PPPoA call with VC-mux encapsulation.
- If the first packet arrives with the LLC encapsulation, then Stinger brings up the PPPoA call with LLC encapsulation.

Verifying the encapsulation type

A standard PPPoA configuration caters to both the types of encapsulation.

To verify the encapsulation type in use, check the `ifmgr` dump of the circuit connection.

The following information displays for a VC encapsulation on executing the `ifmgr` dump:

```
admin >ifmgr -n ckt
Profile[ckt] slot/if=[1/27]
rdh>ifmgr -d 27
iff 0x85077e98
inUse: Yes
hostName: ckt
.
.
.atm1483-encaps : aal5-vc
.
.

admin >info np connection 1 1 27
Connection Handles
if sh/sl act cPort dPort type uctl      udata      dctl      ddata
Encaps: PPPoA over VC - MTU=1524 (50)

Control & Data Service Queues:
-----
ControlQID          324   Lport          0x0000
Scheduler            CBR   CosQueue       1
DataQID             332   Lport          0x0000
Scheduler            UBR   CosQueue
.
.
```

The following information displays for the LLC encapsulation on executing the `ifmgr` dump:

```
admin>ifmgr -n ckt
Profile[ckt] slot/if=[1/27]
rdh>ifmgr -d 27
iff          0x85077e98
```



```

inUse:          Yes
hostName:       ckt
.
.
atm1483-encaps  :      aal5-llc
.
.

admin>info np connection 1 1 27
Connection Handles
if  sh/sl act cPort dPort type  uctl      udata      dctl      ddata
27   1/1  Y  000000 000000 LIM   0/149     0/148     0/149     0/148
Encaps: PPPoA over LLC - MTU=1524 (50)

Control & Data Service Queues:
-----

ControlQID              340  Lport              0x0000
Scheduler                CBR  CosQueue           1


DataQID                  344  Lport              0x0000
Scheduler                UBR  CosQueue           3
.
.

```

Example of configuring a PPPoE connection

Figure 4-25 shows a Stinger IP DSLAM. Across a DSL interface, a CPE bridging device connects to an Ethernet segment with a PPPoE client.

Figure 4-25. Example of a PPPoE session on a DSL interface



Note You can use bidirectional CHAP authentication for PPPoE connections, but it is not required. For information about using bidirectional CHAP, see “Sample PPPoA connection with bidirectional CHAP authentication” on page 4-93.

Overview of PPPoE CONNECTION settings

For background information about IP routing configurations, see “Configuring IPoA subscriber connections” on page 4-49. In addition to those settings, following are relevant PPPoE parameters, shown with default setting, including PAP authentication:

```

[in CONNECTION/"" ]
station* = ""
encapsulation-protocol = atm-circuit
[in CONNECTION/"" :ip-options]
ip-routing-enabled = yes

```

```

remote-address = 0.0.0.0/0
[in CONNECTION/"" :ppp-options]
recv-password = ""
[in CONNECTION/"" :pppoe-options]
pppoe = no
bridge-non-pppoe = no
[in CONNECTION/"" :telco-options]
call-type = ft1

```

Parameter	RADIUS attribute	Setting
station	User-Name (1)	Name of the PPP client system. The value is case sensitive, and must exactly match the name the client presents during authentication.
encapsulation-protocol	Framed-Protocol (7)	Encapsulation protocol. Set to ppp for PPPoE clients.
ip-routing-enabled	Ascend-Route-IP (228)	Enables or disables IP routing for the interface. IP routing is enabled by default.
remote-address	Framed-IP-Address (8) Framed-IP-Netmask (9)	IP address of the remote device, which can include a subnet specification.
recv-password	Password (2)	Password sent by the PPPoE client.
pppoe	Ascend-PPPoE-Enable (74)	Enables or disables processing of PPPoE packets. PPPoE is disabled by default. When both pppoe and pppoa are set to yes in the same CONNECTION profile, the system automatically detects which protocol is used in the incoming traffic stream and performs the required processing. To enable only PPPoE processing, set this parameter to yes and the pppoa-options:pppoa parameter to no .
bridge-non-pppoe	Ascend-Bridge-Non-PPPoE (75)	<i>Not currently supported.</i>
call-type	Ascend-Call-Type (177)	Set call-type to off in PPP profiles. (For RADIUS profiles, set the attribute to Switched .) For background information, see “Recommended call-type setting for PPP sessions” on page 4-85.

Sample PPPoE connection using PAP authentication

First, configure an ATM circuit between the DSLCPE interface and the LIM’s internal interface. See “Terminating traffic on a LIM internal interface” on page 4-89.

Then, verify that the ANSWER-DEFAULTS profile enables the PPP authentication methods to be used and configure a CONNECTION profile to the PPPoE client. For example, the following commands configure a profile using PPP authentication for the PPPoE client in Figure 4-25:

```

admin> new connection pppoe-1
admin> set active = yes

```

```
admin> set encapsulation-protocol = ppp
admin> set ip-options remote-address = 2.2.2.1/29
admin> set ppp-options recv-password = pppoe1!pw
admin> set pppoe-options pppoe = yes
admin> set telco-options call-type = off
admin> write -f
```

Following is a comparable RADIUS user profile:

```
pppoe-1 Password = "pppoe1!pw"
  Service-Type = Framed-User,
  Framed-Protocol = PPP,
  Framed-IP-Address = 2.2.2.1,
  Framed-IP-Netmask = 255.255.255.248,
  Ascend-Call-Type = Switched,
  Ascend-PPPoE-Enable = PPPoE-Yes,
  Ascend-Call-Type = 0
```

Sample PPPoE connection on a VDSL port

On very high speed DSL (VDSL) ports, PPPoE connections use Ethernet QoS on the drop rather than ATM QoS. This example shows configuration of a PPPoE connection on a VDSL port.

The following commands configure an ETHER-QOS profile with priority 0 (the default value, which represents the lowest priority) and a committed information rate (CIR) of 1200Kbps:

```
admin> new ether-qos sats
admin> set committed-information-rate = 1200
admin> write -f
```

For background information about the ETHER-QOS profile for VDSL ports, see the *Stinger VDSL Line Interface Module (LIM) Guide*.

The following commands configure the terminating circuit for the PPPoE connection on VDSL interface 154:

```
admin> new connection atm.ckt.4.1
admin> set active = yes
admin> set encapsulation-protocol = atm-circuit
admin> set ether-options usr-dn-stream-contract = sats
admin> set atm-options vci = 35
admin> set atm-options nailed-group = 154
admin> set atm-options atm1483type = aal5-llc
admin> set atm-connect-options vci = 36
admin> set atm-connect-options nailed-group = 2223
admin> set atm-connect-options atm1483type = aal5-llc
admin> write -f
```

The following commands configure a PPPoE CONNECTION profile. Notice that the CONNECTION profile for PPPoE on VDSL ports uses the same settings as a profile for PPPoE on other types of subscriber ports.

```
admin> new connection sats-pppoe
admin> set active = yes
admin> set encapsulation-protocol = ppp
admin> set ip-options remote-address = 11.11.11.11/8
admin> set ip-options local-address = 11.11.11.111/8
admin> set ip-options multicast-allowed = yes
admin> set ip-options multicast-max-groups = 100
admin> set ppp-options recv-password = ascend
admin> set mp-options enabled = no
admin> set pppoe-options pppoe = yes
admin> write -f
```

Enabling LCP keepalives

RFC 1661, *The Point-to-Point Protocol (PPP)*, describes Link Control Protocol (LCP) Echo messages that may be generated periodically by both ends of a connection, and answered by the opposite end of the connection, to verify connection status. LCP Echo packets generated for this purpose are referred to as *LCP keepalives*.

You can configure the system to generate LCP keepalives, so it can quickly detect and terminate stale sessions when a PPP client has disconnected but the DSL line and ATM connection remains active. The system always replies to LCP keepalives.

RFC 1661 also describes the transmission of magic numbers to detect looped-back links and other data-link layer anomalies, so the system can quickly terminate invalid links. Magic numbers are negotiated when Link Quality Monitoring (LQM) is enabled, but you can also enable magic number negotiation explicitly, so it is performed even if LQM is disabled.

When the Stinger system has been configured to generate LCP keepalives for verifying remote PPP peer status, the system sends an LCP Echo Request and checks for a matching LCP Echo Reply. If the remote PPP peer fails to answer a certain number of consecutive LCP Echo Requests, the PPP connection is terminated.

Parameters and RADIUS attribute-value pairs

Following are the related parameters, shown with default settings:

```
[in ANSWER-DEFAULTS:ppp-answer]
lcp-keepalive-period = 0
lcp-missed-keepalives = 4

[in CONNECTION/"" :ppp-options]
lcp-keepalive-period = 0
lcp-missed-keepalives = 4
```

Parameter	Setting
lcp-keepalive-period	Number of seconds between transmissions of LCP keepalives, from 0 to 65535. With the default zero value, LCP keepalives are disabled. Setting this parameter to a non-zero value enables LCP keepalive. The system responds to all LCP Echo Requests it receives, regardless of this setting.
lcp-missed-keepalives	Number of consecutive unanswered LCP keepalives before the connection is terminated. The valid range is from 1 to 100. When a connection is terminated because the number of unanswered LCP keepalives exceeded the configured maximum, the disconnect code will be set to DIS_PPP_LCP_KEEPAIVE_FAILURE with a value of 172.

These parameters have the following corresponding RADIUS attribute-value pairs:

LCP keepalive parameters

ATTRIBUTE	Ascend-LCP-Keepalive-Period	321	integer
ATTRIBUTE	Ascend-LCP-Keepalive-Missed-Limit	322	integer

Using these attributes requires a server capable of encoding 16-bits VSAs (vendor specific attributes), such as Lucent's Navis Radius. These attribute-value pairs override the corresponding parameter values set in the `answer-defaults` profile.

Sample LCP keepalive configuration

This sample configuration creates a PPP connection profile using LCP keepalives for a remote CellPipe® on a DSL interface. The following commands configure the `answer-defaults` profile to require PPP authentication and to enable LCP keepalives:

```
admin> new answer-defaults
admin> set ppp-answer receive-auth-mode = any-ppp-auth
admin> set ppp-answer lcp-keepalive-period = 4
admin> set ppp-answer lcp-missed-keepalives = 5
admin> write -f
```

PPP connections require two profiles, one for an ATM circuit that terminates on the LIM internal interface, and another for the PPP traffic stream. The following commands configure an ATM circuit that terminates on the LIM internal interface:

```
admin> new connection
admin> set station = atm-ckt-15-1
admin> set active = yes
admin> set telco-options nailed-groups = 701
admin> set atm-options vci = 36
admin> set atm-options nailed-group = 701
admin> set atm-connect-options vci = 36
admin> set atm-connect-options nailed-group = 2797
admin> write -f
```

The next set of commands configures a PPPoE connection to the remote CellPipe®:

```
admin> new connection
admin> set station = pipe
admin> set active = yes
admin> set encapsulation-protocol = ppp
admin> set ip-options remote-address = 10.10.10.1/24
admin> set ip-options local-address = 20.20.20.1/24
admin> set telco-options call-type = off
admin> set ppp-options recv-password = pipe
admin> set ppp-options link-compression = none
admin> set ppp-options lcp-keepalive-period = 4
admin> set ppp-options lcp-missed-keepalives = 5
admin> write -f
```

Following is a comparable RADIUS user profile:

```
pipe    Password = "pipe"
        Framed-Protocol = PPP,
        Framed-Address = 10.10.10.1,
        Framed-Netmask = 255.255.255.0,
        Ascend-PPP-Address = 20.20.20.1,
        Ascend-IF-Netmask = 255.255.255.0,
        Ascend-LCP-Keepalive-Period = 4,
        Ascend-LCP-Keepalive-Missed-Limit = 5
```

LCP keepalive interaction with LQM

Table 4-5 lists system behavior when LCP keepalives are disabled (shown in the top three rows) and when they are enabled (shown in the bottom three rows)

Table 4-5. Configurable LCP Echo interaction with LQM

LCP keepalive enabled?¹	LQM enabled?²	LQM negotiated?³	What is sent or verified
no	no	n/a	Nothing
no ⁴	yes	no	LQM uses LCP Echo
no	yes	yes	LQM reports
yes	no	n/a	LCP keepalives
yes ⁵	yes	no	LCP keepalives
yes ⁶	yes	yes	LCP keepalives and Link-Quality-Reports

1. This column indicates whether LCP keepalives have been enabled by setting the `lcp-keepalive-period` parameter to a non-zero value.

2. This column indicates whether LQM has been enabled by setting the `lqm` parameter to `yes`.

3. This column shows the result of LQM negotiation. A `yes` entry means the PPP peer accepted LQM, and a `no` value means the PPP peer rejected the LQM option.

4. With the conditions shown in this row, the system sends LCP Echoes to gauge the quality of the link. The send rate is controlled via the LQM parameters because this is an LQM feature.
5. With the conditions shown in this row, the PPP peer rejected LQM, so LQM cannot be used on the connection. In this case, the system sends LCP keepalives.
6. With the conditions shown in this row, both LQM reports and LCP keepalives are used to monitor the status of the PPP peer.

Enabling magic-number negotiation when LQM is disabled

You can set the following parameter, shown with its default setting, to enable magic number negotiation even if LQM is disabled.

```
[in ANSWER-DEFAULTS/ppp-answer]  
use-magic-number = no
```

Parameter	Setting
use-magic-number	Enables or disables negotiation of a magic number during LCP negotiation when LQM is not enabled. If this parameter is set to yes and LQM is disabled, the system inspects the magic-number field in received packets and processes a packet normally if the field is equal to zero or the peer's unique magic number. If this parameter is set to no , LQM must be enabled for the system to perform magic-number negotiation.

This setting has no effect if LQM is enabled.

For example, the following commands enable magic-number processing in the **answer-defaults** profile:

```
admin> read answer-defaults  
admin> set ppp-answer use-magic-number = yes  
admin> write -f
```

If LQM is enabled, magic number is negotiated regardless of the setting of this parameter. When LQM is disabled, this setting causes the system to negotiate magic numbers with PPP peers.

Enabling support for PPPoE service names

RFC 2516, *A Method for Transmitting PPP Over Ethernet (PPPoE)*, defines a service-name tag that can be used, for example, to include an ISP name or quality of service indication in PPPoE packets. This section describes how to enable the Stinger system to act as a PPPoE server that filters PPPoE packets based on the service-name field, so that only PPPoE packets that specify certain service names are accepted or rejected.

Details of the service-name negotiation process

PPPoE is negotiated in two phases: a discovery phase and a session phase. The system exchanges Service-Name information during the discovery phase.

- A client initiates a session with a PPPoE Active Discovery Initiation (PADI) packet.

- If the PADI packet specifies an unacceptable service name, the system ignores the packet.
- If the PADI packet specifies an acceptable service name, the system responds with a PPPoE Active Discovery Offer (PADO) packet.
- On receipt of a PADO packet, the client sends a PPPoE Active Discovery Request (PADR) packet.
 - If the PADR packet specifies an unacceptable service name, the system rejects the session request by returning a PPPoE Active Discovery Session-confirmation (PADS) packet with an empty Service-Name-Error tag and a Session ID of 0.
 - If the PADR packet specifies an acceptable service name, the system returns a PADS packet with a valid session ID.

The system examines service-name values at the ATM LLC link level, which is the atm-circuit connection that carries PPPoE traffic for what can be multiple PPPoE sessions.

Configuration overview

Each LLC link carrying PPPoE traffic across a DSL port can apply a PPPoE server configuration that has the following specifications:

- A list of up to 16 service names
- One of the following valid actions to take for PADI and PADR packets that specify one of the service names in the list:
 - Accept (and reply)
 - Ignore (for PADI) / Reject (for PADR)

If a packet specifies a service name that is not in the list, the system takes the opposite action.

- One of the following valid actions to take for PADI or PADR packets that contain an empty Service Name Tag:
 - Accept (and reply)
 - Ignore (for PADI) / Reject (for PADR)

You can create as many PPPoE server configurations as you need (limited only by memory considerations). You can add or modify service names at any time. The changes do not impact established connections, but will impact any new connection that sends PADI and PADR packets.

To simplify configuration, you can define one PPPOE-SERVER profile as the default to be used by all ATM circuits that do not explicitly specify a PPPOE-SERVER configuration.

Creating a PPPOE-SERVER profile

Following are the contents of the PPPOE-SERVER profile, shown with default values:

```
[in PPPOE-SERVER/"""]
name* = ""
empty-service-name-action = accept
service-name-list = [ "" "" "" "" "" "" "" "" "" "" "" "" "" "" "" "" ]
list-service-name-action = ignore
```


The corresponding RADIUS attributes are Ascend vendor-specific, 16-bit. For details about configuring a PPPoE server profile in RADIUS, see “Sample PPPoE server configuration using RADIUS” on page 4-106.

Parameter	RADIUS attribute	Setting
name	Ascend-PPPOE-Server (314)	Name of the profile, up to 23 characters. Two profiles of the same type cannot have the same name.
empty-service-name-action	Ascend-PPPOE-Empty-Service-Name-Action (315)	Action to take when receiving a PADI or PADR with an empty Service Name Tag. accept (0) Reply to the PADI with a PADO that specifies the same (empty) Service Name Tag to accept it. This is the default value, to ensure compatibility with previous system behavior. ignore (1) Ignore the PADI. This means the system will not receive a PADR from the same client with the same service name.
service-name-list	Ascend-PPPOE-Service-Name (316)	A list of up to 16 PPPoE service names. The service-name-list is treated as a whole, so all of the listed names are either honored or ignored, as specified in the list-service-name-action setting.
list-service-name-action	Ascend-PPPOE-Service-Name-Action (317)	Action to take when receiving a PADI or PADR with a service name that is included in the service-name-list. accept (0) Reply to the PADI that contains a service name from the service-name-list with a PADO indicating that the system accepts it. When the client sends a PADR, the Stinger system replies with a PADS that contains a new Session ID to accept the session. With this setting, all non-empty service names that are not included in the service-name-list are rejected.

Parameter	RADIUS attribute	Setting
		ignore (1) Ignore a PADI if it contains a service name from the <code>service-name-list</code>. This is the default value. With the default empty list, the system accepts all service names in a PADI. With this setting and a non-empty <code>service-name-list</code>, the system rejects all service names from the list and accepts all others. If a client sends a different service name in a PADI and PADR, the system checks the name in the PADR and takes the specified action. To reject a PADI, the system sends a PADS with a session ID of 0 and an empty Service-Name-Error tag.

Applying a server profile to a CONNECTION profile

Each `atm-circuit` connection used to terminate PPPoE traffic on the LIM's ATM internal interface can be associated with one PPPoE server profile. This provides the flexibility of selecting the PPPoE server profile to be used at the level of the ATM interface (by the unique combination of DSL port number and VPI-VCI pair).

Following is the relevant parameter, shown with its default setting:

```
[in CONNECTION/"" :pppoe-options]
server-profile = ""
```

The corresponding RADIUS attribute is Ascend vendor-specific, 16-bit.

Parameter	RADIUS attribute	Setting
server-profile	Ascend-PPPOE-Server (314)	Name of a PPPOE-SERVER profile. If this parameter is not configured and the system profile specifies a valid <code>pppoe-default-server-name</code>, that PPPOE-SERVER profile is used. If both this parameter and the SYSTEM profile <code>pppoe-default-server-name</code> parameter are not configured, the internal factory default PPPOE-SERVER profile is used, providing compatibility with previous system behavior. Dependencies The CONNECTION profile must specify <code>atm-circuit</code> encapsulation.

Specifying a default server profile in the SYSTEM profile

Following is the parameter, shown with its default value, for defining a system-wide customized default PPPOE-SERVER profile:

```
[in SYSTEM]
pppoe-default-server-name = default
```

Parameter	Setting
pppoe-default-server-name	Name of the default PPPOE-SERVER profile, to be used whenever a CONNECTION has no server-profile configured, or it has one configured but the specified PPPOE-SERVER profile could not be found. At boot time, if this pppoe-default-server-name has the default value (default) and if no PPPOE-SERVER profiles have been configured locally, the system creates a PPPOE-SERVER profile named default with factory defaults values. This typically happens on the first reset after an upgrade.

You can configure the pppoe-default-server-name setting as a security enhancement, to change the factory default “pass-all” action for CONNECTION profiles that do not assign a PPPOE-SERVER profile. You can also use it to greatly simplify PPPoE service name configuration, if your application requires only system-wide settings. In that case, you can configure only one PPPOE-SERVER profile, specify it as the pppoe-default-server-name value, and leave the CONNECTION profile server-name settings empty.

Specifying a password for RADIUS-defined server profiles

The password to use for all RADIUS-defined PPPoE servers profiles is set by means of the following parameter, shown with its default value:

```
[in EXTERNAL-AUTH:password-profile]
pppoe-servers = ascend
```

Parameter	Setting
pppoe-servers	Password for RADIUS PPPoE server profiles, up to 21 characters.

Sample PPPoE server configuration using local profiles

In the following example, the subscriber PPPoE connection is configured on a 72-port LIM in slot 4 of the system.

- 1 Create a PPPOE-SERVER profile.
admin> new pppoe-server pppoe-srv-gold
admin> set service-name-list 1 = data
admin> set service-name-list 2 = video
admin> set list-service-name-action = accept
admin> write -f
- 2 Determine the nailed-group number of port 1 of the LIM in slot 4:
admin> which -n { 1 4 1 }
Nailed group corresponding to port { shelf-1 slot-4 1 } is 151
- 3 Determine the nailed-group of the internal interface of a 72-port LIM in slot 4:
admin> which -n { 1 4 73 }

- Nailed group corresponding to port { shelf-1 slot-4 73 } is 2271
- 4 Enable SDSL interface 1 in slot 4.
admin> read sdsl { 1 4 1 }
admin> set enabled = yes
admin> write -f
 - 5 Configure an ATM circuit from the DSL interface to the LIM's ATM internal interface:
admin> new connection cir-1
admin> set active = yes
admin> set encapsulation-protocol = atm-circuit
admin> set atm-options vci = 38
admin> set atm-options nailed-group = 151
admin> set atm-options atm1483type = aal5-llc
admin> set atm-connect-options vci = 36
admin> set atm-connect-options nailed-group = 2271
admin> set atm-connect-options atm1483type = aal5-llc
admin> set pppoe-options server-profile = pppoe-srv-gold
admin> write -f
 - 6 Configure a PPP CONNECTION profile for the PPPoE subscriber.
admin> new connection pppoe-1
admin> set active = yes
admin> set encapsulation-protocol = ppp
admin> set ip-options remote-address = 2.2.2.1/29
admin> set ppp-options recv-password = pppoe1!pw
admin> set pppoe-options pppoe = yes
admin> set telco-options call-type = off
admin> write -f

Sample PPPoE server configuration using RADIUS

Following is a sample RADIUS PPPoE server profile:

```
permpppoe-stg123-1 Password= "ascend"  
User-Service = Dialout-Framed-User,  
Ascend-PPPOE-Server = "pppoe-remote",  
Ascend-PPPOE-Empty-Service-Name-Action = accept,  
Ascend-PPPOE-Service-Name = "video",  
Ascend-PPPOE-Service-Name = "internet",  
Ascend-PPPOE-Service-Name-Action = accept
```

Table 4-6. Permanent RADIUS PPPoE server profiles

Element	Description
permpppoe-stg123-1	The user name for RADIUS PPPoE server profiles must begin with permpppoe- followed by the system name if one has been defined, followed by a hyphen. For example, if the system name is stg123, all PPPoE server profiles must begin with the string permpppoe-stg123-. If the system name has not been defined, all PPPoE server profiles must begin with the string permpppoe--. The last element in the PPPoE server profile user name is the profile index, which must start with 1 and monotonically increase by 1, with no gaps in the sequence, for example: <pre>permpppoe-stg123-1 permpppoe-stg123-2 permpppoe-stg123-3</pre> A gap in the sequence of index numbers would cause the system to ignore subsequent profiles in the list.
Password (2)	The password to use for all RADIUS-defined PPPoE servers profiles is set by means of the following parameter, shown with its default value: <pre>[in EXTERNAL-AUTH:password-profile] pppoe-servers = ascend</pre>
User-Service (6)	For PPPoE server profiles, this attribute-value pair must specify Dialout-Framed-User.
Ascend-PPPOE-Server (314)	Name of the system-generated RADIUS pppoe-server profile that corresponds to this pseudo-user profile. The profile name is visible (read only) when you issue the following command: <pre>admin> dir pppoe-server</pre> If this name is not specified in the pseudo-user profile, the system ignores the profile and logs an error message.
Ascend-PPPOE-Empty-Service-Name-Action (315)	Action to take when receiving a PADI or PADR with an empty Service Name Tag. 0 (accept) and 1 (ignore) are valid values. For details, see the description under “Creating a PPPOE-SERVER profile” on page 4-102.
Ascend-PPPOE-Service-Name (316)	Add one of these attribute-value pairs for each service name, up to 16 names. (Any additional service names beyond 16 are ignored.)

Table 4-6. Permanent RADIUS PPPoE server profiles (Continued)

Element	Description
Ascend-PPPOE-Service-Name-Action (317)	Action to take when receiving a PADI or PADR with any service name that is listed in an Ascend-PPPOE-Service-Name attribute-value pair in this profile. 0 (accept) or 1 (ignore) are valid values. Like in the local profile, the Ascend-PPPOE-Service-Name-Action applies to all Ascend-PPPOE-Service-Name attributes. For details, see the description under “Creating a PPPOE-SERVER profile” on page 4-102.

If you modify the profile during run-time

When you change the parameters of a RADIUS PPPoE server profile and then issue a `refresh -e` command, changes take effect for the next PADI or PADR packet on any valid connection to which the PPPoE server profile is linked.

If you configure an invalid action

If Ascend-PPPOE-Server specifies a valid name, the system accepts the pseudo-user profile and creates a system-generated PPPOE-SERVER profile, even if you have configured an improper value for the Ascend-PPPOE-Empty-Service-Name-Action or Ascend-PPPOE-Service-Name-Action attribute. In that case, the system uses the factory default value rather and ignores the invalid value. This is consistent with system behavior when you restore a configuration.

When RADIUS PPPoE server profiles are retrieved from RADIUS

PPPoE server pseudo-users are retrieved from RADIUS at the following times:

- At boot time when the first IP interface comes up
- When you issue a `refresh -a` or `refresh -e` command
The `-e` option causes the system to retrieve all PPPoE server profiles (and only those profiles) from the RADIUS server. As with previous software versions, the `refresh -a` command causes the system to retrieve all permanent RADIUS profiles, including PPPoE server profiles.
- When you send an SNMP SET command on the `sysConfigRadiusCmd` object from `ascend.mib` with a value of 1 (all) or 7 (pppoe).
- When you make a significant change to the EXTERNAL-AUTH profile, such as enabling RADIUS, adding, removing, or changing a RADIUS authentication server hosts list, changing the RADIUS authentication ports, and so forth.

When RADIUS-learned PPPoE server profiles are deleted

The system deletes profiles it has learned from RADIUS under the following conditions:

- If you disable RADIUS in the EXTERNAL-AUTH profile, all RADIUS-learned profiles are deleted.
- If you delete a PPPoE server pseudo-user from the RADIUS server, that profile will also be deleted from the Stinger system on the next refresh.

- If you issue a `refresh` command and it fails due to an invalid password, all RADIUS-learned profiles are deleted and no retry is allowed. (If the command fails due to a time-out or network error, the system retries and profiles are not affected.)

Syslog messages related to RADIUS PPPoE server configurations

When PPPoE servers have been successfully refreshed, you will see the following message of level INFO:

```
"Radius permpppoe Profiles refreshed"
```

If you successfully retrieve a PPPoE server profile from RADIUS, but it has invalid parameters, then you will see the following message of level ERROR:

```
"Received invalid RADIUS PPPOE-SERVER profile (<detail>)"
```

where `<detail>` is a custom message providing more information on what is invalid. This message is exactly the same message that you would have seen if you had tried to write the invalid profile locally through the CLI. For example:

```
LOG error, Shelf 1, Controller-1, Time: 19:48:37--  
Received invalid RADIUS PPPOE-SERVER profile ("name" must be specified.)
```

The above message indicates that an Ascend-PPPOE-Server attribute with an empty string was received. The Ascend-PPPOE-Server attribute corresponds to the `name` parameter of a PPPOE-SERVER profile.

If you successfully retrieve a PPPoE server profile from RADIUS, but it has a name conflicting with a local `pppoe-server` profile, you will see the following message of level ERROR:

```
"Could not add duplicate RADIUS PPPOE-SERVER profile' <profile-name>'"
```

where `<profile-name>` is the name of the local profile we could not overwrite. For example:

```
LOG error, Shelf 1, Controller-1, Time: 19:58:51--  
Could not add duplicate RADIUS PPPOE-SERVER profile 'pppoe-video'
```

Administrative tools for IP routing

The system supports several commands that are useful for locating the sources of problems on an IP network and for communicating with other hosts for management purposes. For examples, see the chapter on working with IP traffic in the *Stinger Administration Guide*, and entries in the *Stinger Reference* on commands such as the following:

- `arptable`
- `ipcache`
- `iproute`
- `netstat`
- `nslookup`
- `ping`
- `telnet`
- `traceroute`

Ethernet and IP QoS


5

Overview of the QoS implementation	5-1
Introduction to PACKET-FLOWS profile settings	5-5
QoS-related connection and interface settings.	5-22
QoS-related settings in the SYSTEM profile.	5-25
Examples of configuring QoS.	5-28
Administrative tools for monitoring IP QoS.	5-47
Limitations with the current software version.	5-61

Overview of the QoS implementation

To provide quality of service (QoS), the system classifies an incoming packet flow according to configured rules, rate-limits the traffic, assigns packets a priority, and finally marks the outgoing traffic. Figure 5-1 shows the order in which these processes occur.

Figure 5-1. QoS subsystems



This scheme provides QoS in one direction of traffic flow only. Rate limiting, scheduling, and packet marking occur after the packets have been classified, and are always applied at the output side of the outgoing interface.

Table 5-1. Descriptions of QoS support for routed, bridged, and VLAN traffic

Type of QoS	Description
IP QoS	IP QoS uses packet classifiers to select packets. IP QoS supports: ■ priority scheduling with rate limiting ■ IP ToS marking IP QoS can be applied to all terminated routed connections, as well as to routed VLAN interfaces. For routed VLAN interfaces, both IP ToS marking and Ethernet priority marking are supported.
Ethernet QoS	Ethernet QoS can use layer 2 or packet classifiers (but not both) to select packets. Ethernet QoS supports: ■ strict priority scheduling with rate limiting ■ Ethernet priority marking Ethernet QoS can be applied to all bridged connections and to bridged and routed VLAN Ethernet interfaces. For routed VLAN interfaces, both IP ToS marking and Ethernet priority marking are supported.

Packet classification subsystem

The classification subsystem selects packets based on the content of packet headers according to defined rules. Rate limiting, priority scheduling, and marking can then be performed on selected packets.

A classification rule set can be created with only one rule or with a combination of rules. When a combination of rules is used, the subsystem performs a logical AND between the rules. For example, if the rule set is composed of a source IP address and its netmask and the IP protocol, only packets that match both the specified source address and the specified protocol are selected. In addition, when multiple rules are specified, the rules are applied in the order in which they appear in the profile. The first rule is applied first. If it does not match, then the next one is applied, and so on.



Caution Ethernet fields in a `layer2-classifier` subprofile and IP-specific fields in a `packet-classifier` subprofile *must not* be configured together in the same `PACKET-FLOWS` profile. Configuring both classifications in the same profile yields unpredictable results. Note that the system does not prevent you from such a misconfiguration.

Interface grouping provides a simple way to classify packet flows by input interface. You configure input interfaces with an interface group number, and classify traffic based on that group number alone, or that number in combination with other rules. Interface grouping is configured in the `packet-classifier` subprofile but it is not IP-specific. It can be combined with layer 2 classification in the same `PACKET-FLOWS` profile.

Rate-limiting, prioritization, and scheduling subsystem

Once the classification has been done, the classification result is passed to the scheduling subsystem, which is responsible for selecting a queue according to the priorities, the rate-limiting result, and scheduling at the connection level.

For WAN connections, traffic scheduling and shaping can be influenced by ATM service contracts. For Ethernet IP or VLAN Ethernet interfaces, traffic scheduling and shaping can be explicitly configured in the interface profiles. For details, see “QoS-related connection and interface settings” on page 5-22.

For individual packet flows, up to eight flow priorities (from 0 to 7) can be defined per connection in a PACKET-FLOWS profile. Each priority level has a queue. The scheduler selects the proper output queue according to the priority configured for the packet flow and the rate limiting associated with the traffic class, and the bandwidth of the connection.

The scheduler attached to the connection selects a priority queue as follows:

- 1 Select the highest priority queue.
- 2 If the selected queue has no traffic, select the next highest priority queue.
- 3 If the selected queue has no traffic, select the next highest priority queue, and so forth.

If you are using distributed multicast with high-bandwidth LIMs, see “Changes in packet-flow prioritization with distributed multicast” on page 9-29 for related information.

With rate-limiting, some packets are assigned lower priority queues than the configured priority when bursty conditions prevail. The rate limiting subsystem implements an RFC 2697 Single Rate Three Color Marker to meter an IP packet flow and color-mark its packets either green, yellow, or red. Marking is based on bandwidth allowed to the packet flow, and how far (or how long) the traffic stream exceeds the allowed bandwidth.

The policing algorithm (one rate, three colors) defined by the RFC 2697 is fully supported by the Stinger QoS implementation. Once a flow has been marked green, yellow or red, it is possible to either drop certain colors of packets (such as those marked red) or to reprioritize the flow with different, decreasing levels of priority for yellow or red packets. For example, the system might discard all red packets, forward yellow packets as best effort, and forward green packets with a high priority.

The rate-limiting subsystem always operates in color-blind mode. That is, the metering process assumes that the packet stream it receives has not already been color marked.

Packet marking engine

The packet marking engine marks outgoing packets to receive a particular service level along their path. Table 5-2 shows the types of packet marking supported on various types of interfaces. The *Profile* column in Table 5-2 indicates the profile in which the PACKET-FLOWS profile is applied.

Table 5-2. Packet marking supported on egress interfaces

Profile	Type of egress interface	Marking action
IP-INTERFACE	Routed Ethernet	IP ToS
IP-INTERFACE	Routed VLAN	IP ToS
VLAN-ETHERNET	Bridged VLAN	Ethernet p-bit
CONNECTION	Routed WAN	IP ToS
CONNECTION	Pretagged stacked VLAN	Ethernet p-bit
CONNECTION	Bridged VLAN (other than pretagged) ¹	No marking support

1. Pretagged stacked VLAN connections are bridged, but because these connections contain a VLAN tag when received from the subscriber side, they can support p-bit marking in the downstream direction.



Note Only one type of marking can be performed for all packet flows defined in a given PACKET-FLOWS profile. Both the fields cannot be marked simultaneously.

The system does not support classification of routed packets on the basis of layer 2 parameters, or marking of IP ToS values for bridged packets.

Configuration steps

The configuration of QoS for a particular output interface should be done in the following order:

- 1 Define the packet flows that must be discriminated for the output interface, and specify the scheduling priority for each flow. This step requires the creation of one PACKET-FLOWS profile.
- 2 If needed, group input interfaces and specify the interface group number in the PACKET-FLOWS profile.
- 3 Attach the PACKET-FLOWS profile to the output interface. You do this in the IP-INTERFACE, VLAN-ETHERNET, or CONNECTION profile for the interface. You can also attach the profile to a RADIUS-authenticated connection by using the Ascend-Packet-Flows (324) vendor-specific attribute.

What the system does at the output interface

Once the output interface is up, the system performs following steps:

- 1 Retrieve the PACKET-FLOWS profile attached to the interface
- 2 Perform checks to ensure that the configuration is valid.
- 3 Create the number of queues needed according to the number of priorities defined by the PACKET-FLOWS profile.
- 4 For ATM connections, select the correct scheduler according to the connection-level ATM service contract configured in the ATM-QOS profile. (See “ATM QoS and IP QoS considerations” on page 5-24.)

- 5 For Ethernet IP or VLAN Ethernet interfaces, select the correct scheduler in the IP-INTERFACE or vlan-ethernet priority configuration. (See “Configuring Ethernet IP traffic shaping” on page 5-24.)
- 6 Create the classification rules and apply them to the output interface.

Because classifiers must be communicated to the underlying system hardware, which can be time intensive when many rules have been defined, a method is provided for batching updates system wide. For details, see “Performance recommendations” on page 5-25.

Default IP QoS configuration

When no PACKET-FLOWS profile has been attached to an interface, the following default IP QoS is applied:

- Control traffic goes in the highest priority queue.
- If multicast is enabled on the interface, multicast traffic goes in a high priority queue.
- Unicast traffic goes in a low priority queue.

When a PACKET-FLOWS profile has been attached to the interface but does not contain an explicit default rule, the system creates one default rule for which the scheduling priority will be the lowest priority.

If you are using distributed multicast with high-bandwidth LIMs, see “Changes in packet-flow prioritization with distributed multicast” on page 9-29 for related information.

Introduction to PACKET-FLOWS profile settings

A PACKET-FLOWS profile describes a set of up to 32 packet flows and how each flow should be scheduled and marked. Following are the parameters and subprofiles, shown with default values, at the top level of the profile:

```
[in PACKET-FLOWS/""]  
name* = ""  
marking-type = none  
flow = [ { { 0 0 00:00 } { 0 0.0.0.0/0 0.0.0.0/0 0 0 none 0 none 0 } { 0 } } +  
global-packet-marking = [ { 0 0 0 } { 1 0 0 } { 2 0 0 } { 3 0 0 } { 4 0 0 } {+  
exception-packet-marking = { 8 0 0 }
```

Parameter or subprofile	Purpose
name	Name of the profile, up to 31 characters, used to uniquely identify the packet flow set within the Stinger system.
marking-type	Specify whether to mark an IP ToS value (layer 3) or VLAN p-bit value (layer 2) in the packet headers.
flow	An array of 32 indexed subprofiles for configuring classification rules and a scheduling priority for up to 32 different flows. Each indexed subprofile contains the following subprofiles:

Parameter or subprofile	Purpose
layer2-classifier	Classify an Ethernet bridged traffic flow. A given PACKET-FLOWS profile can configure either layer2-classifier settings or IP-specific packet-classifier settings, but not both.
packet-classifier	Classify an IP packet flow. The packets can be routed through the system or bridged Ethernet traffic. A given PACKET-FLOWS profile can configure either layer2-classifier settings or IP-specific packet-classifier settings, but not both.
scheduling	Specify a priority for the traffic class. For VLAN, strict priority queueing is supported.
global-packet-marking	Specifies the IP ToS or VLAN p-bit values to be marked in the outgoing traffic.
exception-packet-marking	Specifies the IP ToS or VLAN p-bit values for outgoing packets generated by the Stinger system, such as ARP or ICMP packets.

Layer 2 classifiers

Layer 2 classification uses Ethernet header fields to select frames for rate limiting, prioritization, and scheduling, in both upstream and downstream directions.



Note Layer 2 classifiers are used only for Ethernet QoS. The system does not support classification of routed packets on the basis of layer 2 parameters, or the classification of packet flows using a mix of layer 2 and IP-specific packet classification fields.

Following are the parameters, shown with default values, for classifying traffic at layer 2. For each parameter, the default zero or null setting matches *any* value in the corresponding fields of the packet.

```
[in PACKET-FLOWS/"" :flow[1]:layer2-classifier]
vlan-id = 0
ethernet-priority = 0
ethernet-type = 00:00
```

Parameter	Description
vlan-id	The 802.1Q/P standard specifies a tag that appends to a MAC frame. The VLAN tag carries both the VLAN ID (12-bit) and Prioritization (3-bit). This parameter can specify a VLAN ID (a number from 0 to 4095). In the case of stacked VLANs, the user-level <code>vlan-id</code> is used instead of the NSP-level <code>vlan-id</code>. See “Overview of VLAN stacking settings for tagged frames” on page 3-31.
ethernet-priority	802.1Q/P priority value (p-bit value) of the frame’s VLAN tag. A number from 0 to 7. In the case of stacked VLANs, the user-level <code>ethernet-priority</code> is used instead of the NSP-level <code>ethernet-priority</code>. See “Overview of VLAN stacking settings for tagged frames” on page 3-31.
ethernet-type	Ethernet type. A two-byte hexadecimal number representing the Ethernet protocol. For example, PPPoE values are 8863 (PPPoE Discovery Stage packets), and 8864 (PPP Session Stage packets).

For an example of layer 2 classification, see “Configuring bridging VLAN Ethernet QoS” on page 5-36.

Packet classifiers

Packet classification uses an interface group number and/or IP packet header fields to select packets for rate-limiting, prioritization, and scheduling, in both upstream and downstream directions. The IP packets can be routed or bridged.

Following are the parameters, shown with default values, for classifying traffic at IP packet layers. For each parameter, the default zero or null setting matches *any* value in the corresponding fields of the packet.

```
[in PACKET-FLOWS/"" :flow[1]:packet-classifier]
interface-group = 0
source-ip-address = 0.0.0.0/0
source-ip-netmask = 0.0.0.0
destination-ip-address = 0.0.0.0/0
destination-ip-netmask = 0.0.0.0
ip-protocol = 0
ip-tos = 0
source-port-comparison = none
source-port = 0
destination-port-comparison = none
destination-port = 0
```

Parameter	Description
interface-group	Interface group number (from 0 to 255) used for classification. Value 0 matches all interfaces, so it means no classification will be performed based on interface group. Up to 256 groups can be created. Interface grouping is configured in the <code>packet-classifier</code> subprofile but it is not IP-specific. It can be combined with layer 2 classification in the same PACKET-FLOWS profile.
source-ip-address	This parameter and the optional <code>source-ip-netmask</code> are used for a classification based on the source IP address of the incoming packets. For example, 192.207.23.13/29.
source-ip-netmask	Network mask to apply before the source IP address comparison. For example, 255.255.255.248.
destination-ip-address	This parameter and the optional <code>destination-ip-netmask</code> are used for a classification based on the destination IP address of the incoming packets. For example, 192.207.23.13/29.
destination-ip-netmask	Network mask to apply before the destination IP address comparison. For example, 255.255.255.248.
ip-protocol	Numeric value (from 0 to 255), used for a classification based on the IP protocol. For a list of assigned protocol numbers, see RFC 1700, <i>Assigned Numbers</i> .
ip-tos	Numeric value (from 0 to 255), used for a classification based on the IP TOS bits (8 bits) present in the incoming packets.
source-port-comparison	Type of comparison to be done on the source UDP/TCP port for a classification based on source port. Enumerated field, values: none Packet address/port is not used in classifying this flow. less Packet address/port less than rule's address/port. equal Packet address/port same as rule's address/port. greater Packet address/port greater than rule's address/port.
source-port	Numeric value (from 0 to 65535) used for the source UDP/TCP port comparison defined by <code>source-port-comparison</code> .
destination-port-comparison	Type of comparison to be done on the destination UDP/TCP port for a classification based on destination port. Enumerated field, values: none Packet address/port is not used in classifying this flow.

Parameter	Description
	less Packet address/port less than rule's address/port.
	equal Packet address/port same as rule's address/port.
	greater Packet address/port greater than rule's address/port.
destination-port	Numeric value (from 0 to 65535) used for the destination UDP/TCP port comparison defined by destination-port-comparison.

Caveat about fragmented IP packets

Keep in mind that for fragmented IP packets, only the IP packet header can be used for classification, which means that classification based on UDP or TCP source and destination ports cannot be used.

Details of packet classifier comparison passes

In evaluating incoming traffic destined for an particular output interface, the system begins by comparing the packet header to rule #1. If that comparison does not result in a match, it proceeds to compare the packet to rule #2, and so forth. The comparison process stops immediately if conditions in a packet match a classification rule. At that point, the scheduling and marking in the profile are applied to the packet flow. Within each rule, the comparisons proceed as shown in Table 5-3:

Table 5-3. Comparison passes performed on inbound packet flows

Comparison pass	Description
1. Interface group	If the group number assigned to the incoming interface does not match the <code>interface-group</code> of the rule, the comparison fails. See “Prioritizing IP packet flows based on DSL service contracts” on page 5-29 for an example.
2. Source address/netmask	If the source address of the packet does not match the <code>source-address/source-netmask</code> of the rule, the comparison fails.
3. Destination address/netmask	If the destination address of the packet does not match the <code>destination-address/destination-netmask</code> of the rule, the comparison fails.
4. IP protocol	If the protocol field of the packet does not match the <code>ip-protocol</code> of the rule, the rule comparison fails.
5. IP TOS bits	If the 8 TOS bits in the incoming packet header do not match the rule's <code>ip-tos</code> , the comparison fails.
6. UDP/TCP source port	If the source address of the packet does not compare as specified to the <code>source-port</code> of the rule, the comparison fails.

Table 5-3. Comparison passes performed on inbound packet flows (Continued)

Comparison pass	Description
7. UDP/TCP destination port	If the destination address of the packet does not compare as specified to the destination-port of the rule, the comparison fails.

For each comparison, the default zero or null setting for an identifier matches all values in a packet. If all comparisons succeed, the packet matches the rule. If any one of the comparisons fail, the packet does not match the rule.

Comparisons of IP addresses

When you specify a source or destination address in a rule, the system assigns a priority to packets received from or sent to that address. If you also specify a subnet mask, the system applies the mask to the address value before comparing the resulting value to the source or destination address in a packet.

To apply the mask, the system translates both the mask and address values into binary format and then uses a logical AND to apply the mask to the address. The mask hides the bits whose positions match those of the binary zeros in the mask. A mask of all zeros (the default) masks all bits. If the address value itself is also all zeros (the default), the rule matches any source or destination address. A mask of all ones (255.255.255.255) masks no bits, so the full address for a single host is compared to the address value.

You can use the address mask to mask out the host portion of an address or the host and subnet portion, so the specification matches the address to or from any host on a given network. For example, the following sample commands specify that any packet received from subnet 192.168.2.0 must be processed at priority level 3:

```
admin> read packet-flows src-netmask-demo
admin> set flow 1 packet-classifier source-ip-netmask = 255.255.255.0
admin> set flow 1 packet-classifier source-ip-address = 192.168.2.2
admin> set flow 1 scheduling queue-priority = 3
admin> write -f
```

Comparisons of IP TOS values

A rule specifying an ip-tos value allows for prioritization of traffic based on the value of the IP TOS byte in the IPv4 header. The default zero value matches all values in the TOS byte. A nonzero value must be matched exactly. For example, the following sample commands specify that packets with a TOS value of 255 must be processed at priority level 5:

```
admin> read packet-flows tos-demo
admin> set flow 1 packet-classifier ip-tos = 255
admin> set flow 1 scheduling queue-priority = 5
admin> write -f
```

Comparisons of port numbers

Rules can specify a port number to be compared to the source or destination port (or both) in a packet.



Note When a nonzero source-port or destination-port is specified in a PACKET-FLOWS profile, the `ip-protocol` value must be set to either 17 (UDP) or 6 (TCP) and the `source-port-comparison` or `destination-port-comparison` parameter must specify the type of comparison to be performed.

TCP and UDP port numbers are typically assigned to services. For a list of well-known port assignments, see RFC 1700, *Assigned Numbers*.

The `source-port-comparison` or `destination-port-comparison` setting determines when a match occurs. For all traffic, the following settings can be used to compare source port or destination port values:

- `none` (no comparison is made)
- `equal` (equal to)
- `less` (less than)
- `greater` (greater than)

The following commands specify that packets received on any port other than port 69 (which is used for TFTP) should be assigned priority level 2:

```
admin> read packet-flows src-port-demo
admin> set flow 1 packet-classifier ip-protocol = 17
admin> set flow 1 packet-classifier source-port-comparison = less
admin> set flow 1 packet-classifier source-port = 69
admin> set flow 2 packet-classifier ip-protocol = 17
admin> set flow 2 packet-classifier source-port-comparison = greater
admin> set flow 2 packet-classifier source-port = 69
admin> set flow 1 scheduling queue-priority = 2
admin> set flow 2 scheduling queue-priority = 2
admin> write -f
```

How nonmatching packets are prioritized (the default rule)

The default rule, which defines how non-matching packets are prioritized, has following characteristics:

- The default rule must be the last `flow` configuration in the PACKET-FLOWS profile.
- All `packet-classifier` fields and `layer2-classifier` of the default rule must be 0 or `none` (the default values).
- The `scheduling queue-priority` for the default rule must be set to the desired priority setting for nonmatching packets.

If you do not define a scheduling priority for the default rule, the system uses priority zero for nonmatching packets. If a rule specifies more than one default rule, only the first one is taken into account. If a PACKET-FLOWS profile specifies 32 classifications and does not include a default rule, the system creates an implicit default rule.

For example, the following commands specify priority 1 for packets that do not match the configured rule for flow 1:

```
admin> read packet-flows src-netmask-demo
admin> set flow 1 packet-classifier source-ip-netmask = 255.255.255.0
admin> set flow 1 packet-classifier source-ip-address = 192.168.2.2
admin> set flow 1 scheduling queue-priority = 3
admin> set flow 2 scheduling queue-priority = 1
admin> write -f
```

Scheduling and rate limiting

The scheduling and rate-limiting settings apply to routed outgoing interfaces and VLAN Ethernet outgoing interfaces. Table 5-4 defines terms used in this section:

Table 5-4. Rate limiting terminology

Term	Definition
CIR	Committed Information Rate
CBS	Committed Burst Size
EBS	Excess Burst Size
Green	Traffic color returned by the policer when the traffic does not exceed CBS.
Yellow	Traffic color returned by the policer when the traffic exceeds CBS but does not exceed EBS.
Red	Traffic color returned by the policer when the traffic exceeds CBS and EBS.

Following are the parameters, shown with default settings, for configuring scheduling and rate limiting:

```
[in PACKET-FLOWS/"" :flow[1]:scheduling]
queue-priority = 0
yellow-queue-priority = -1
red-queue-priority = -1
committed-information-rate = 0
committed-burst-size = 0
excess-burst-size = 0
```

Parameter	Setting
queue-priority	Queue priority of the green traffic (from 0 to 7). The default value of 0 means the lowest priority queue. For scheduling Ethernet traffic, this setting represents a strict priority for the bridged traffic flow.

Parameter	Setting
yellow-queue-priority	Priority (from -1 to 7) of the flow when the policer is active and detects traffic above CBS but below EBS. The default value of -1 indicates that yellow traffic must be discarded. In this condition, the red-priority queue is not used and must be also configured with the value -1.
red-queue-priority	Priority (from -1 to 7) of the flow when the policer is active and detects traffic above EBS. The default value of -1 indicates that red traffic must be discarded.
committed-information-rate	Average bandwidth of the flow in Kbps, from 0 to 1,000,000. With the default value of 0, the policer is disabled and the flow is not rate limited. This is the maximum possible setting for CIR. The flow rate in this case is the rate defined by the ATM-QOS profile for ATM connections, or by the scheduling subprofile for routed (IP-INTERFACE) or VLAN (VLAN-ETHERNET) interfaces. A nonzero value automatically enables the policer in color-blind mode. ■ The minimum policing rate (minimum CIR) with an IP2000 module is 16.276 Kbps rounded to 17 Kbps. ■ With an IP2100 control module, the minimum policing rate (minimum CIR) is 31Kbps. The precision of the policing algorithm is a function of the CIR, the configured peak rate of the connection or interface, and the actual line rate. The worst precision occurs when desired peak rate is close to the line rate.
committed-burst-size	CBS expressed in bytes. This value (from 0 to 65,535) is used to size the CBS token bucket. The minimum recommended value when the policer is enabled is two times the MTU of the connection.
excess-burst-size	EBS expressed in bytes. This value (from 0 to 65,535) is used to size the EBS token bucket. The minimum recommended value when the policer is enabled is three times the MTU of the connection.

Token buckets in the single-rate three color policing algorithm

A “token bucket” is a metaphor for a calculated transfer rate used internally by the policer to regulate a packet flow. In the metaphor, the system places tokens into a bucket at a certain rate (the CIR). The size of a token is the IP packet size, as defined by RFC 2697. If the transfer rate is always at the CIR level, tokens pass through the buckets without accumulating.

When a source sends bursty traffic, the system discards packets if tokens are not available in either of the commit and excess buckets. The largest burst a source can send into the network is generally proportional to the size of the bucket.

It is strongly recommended to configure the size of the buckets to be able to store at least two packets. That is the reason why the minimum recommended value for CBS is two times the MTU of the connection. Because the EBS bucket must be larger than the CBS bucket, the minimum recommended value for EBS is three times the MTU (to store three packets).

Using a single rate two-color algorithm

To configure the rate-limiting subsystem to use only two colors (green and red), set the `excess-burst-size` parameter to 0 and the `yellow-queue-priority` parameter to -1 (not used). You can configure the other parameters in the `scheduling` subprofile to any desired values.

Notes on the policing implementation

- The policer acts as a real policer, so it has no associated queues. This means that rate-limiting a bursty flow implies a correct setting of the burst-size parameters.
- Since a policer is in use, multicast flows cannot be rate-limited. If the multicast group was limited, in fact all member of the group will be rate-limited.
- The minimum policing rate is 17Kbps regardless of the type of interface.
- The policing precision is about 1% in the worst case (when the peak rate is close to the line rate).
- The Stinger supports up to 32K policing flows.
- Some reserves have to be done regarding the throughput of the network processor due to the policer use.
- For fragmented packets, the policer might not be applicable if the flow to be policed includes the IP protocol (TCP/UDP) and associated port number.

Rate limiting on VLAN bridged interfaces

Rate-limiting is performed at the outgoing interface for VLAN bridged traffic as shown in Table 5-5:

Table 5-5. Flow-based rate limiting configurations

Type of interface	Rate limiting description
1:1 VLAN bridging (Ethernet side)	Each connection can have an independent rate-limiting policy for the traffic transmitted over its VLAN Ethernet interface.
1:N VLAN bridging (Ethernet side)	The rate-limiting policy can be applied for the aggregate traffic from multiple user connections being transmitted over the VLAN Ethernet interface.
Stacked VLANs (Ethernet side)	The rate-limiting policy in the upstream direction can be configured on a per NSP VLAN basis, so it would be applicable for the aggregate traffic being transmitted over the NSP VLAN interface.
All types of bridged DSL interfaces	The rate-limiting policy in the downstream direction can be configured independently on a per-connection basis.

Within the single-rate tricolor marker algorithm, the size of the bridged packet is considered as that of the payload, excluding any layer 2 headers or trailers applicable at the IP control module. For example:

- For a bridged-LLC packet received from a transparent-bridged DSL interface, the LLC and MAC header (Destination Address, Source Address and EtherType fields) are excluded while calculating the packet size for rate-limiting purposes.
- For a VLAN-tagged packet received from a transparent-bridged VLAN Ethernet interface, the MAC header, VLAN tag and the Ethernet CRC are excluded while calculating the packet size for rate-limiting purposes.



Note For transparent bridged (1:N VLAN) connections, the rate-limiting parameters are not applicable to flooded traffic. Also, multicast traffic cannot be rate-limited.

Example of rate limiting on a BIR connection

This sample configuration creates a PACKET-FLOWS profile that rate limits and prioritizes a mix of unicast and multicast traffic flows, and applies it to a bridged IP routing (BIR) connection.

- 1 Configure the PACKET-FLOWS profile to assign a high priority (7) to multicast traffic, middle priority (5) to ICMP packets, and low priority (3) to IP unicast traffic. Unicast traffic and ICMP traffic is rate limited.

```
admin> new packet-flows bir-flow
admin> set flow 1 packet-classifier destination-ip-address = 224.1.1.0
admin> set flow 1 packet-classifier destination-ip-net = 255.255.255.0
admin> set flow 1 scheduling queue-priority = 7
admin> set flow 2 packet-classifier ip-protocol = 1
admin> set flow 2 scheduling queue-priority = 5
admin> set flow 2 scheduling yellow-queue-priority = 4
admin> set flow 2 scheduling committed-information-rate = 100
admin> set flow 2 scheduling committed-burst-size = 3000
admin> set flow 2 scheduling excess-burst-size = 30000
admin> set flow 3 packet-classifier ip-protocol = 4
admin> set flow 3 scheduling queue-priority = 3
admin> set flow 3 scheduling committed-information-rate = 200
admin> set flow 3 scheduling committed-burst-size = 3000
admin> write -f
```
- 2 Configure the BIR CONNECTION profile and apply the PACKET-FLOWS profile.

```
admin> new connection bir-user
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 192.168.190.2/24
admin> set ip-options local-address = 192.168.190.1/24
admin> set ip-options multicast-allowed = yes
admin> set ip-options multicast-max-groups = 3
```

```
admin> set ip-options packet-flows = bir-flow
admin> set bir-options enable = yes
admin> write -f
```

Example of rate limiting on a 1:1 bridged VLAN

This sample configuration creates a PACKET-FLOWS profile that rate limits and prioritizes traffic flows, and applies it to the downstream DLS interface of a 1:1 VLAN circuit. This example assumes that the VLAN circuit has already been configured. For background information, see Chapter 3, “VLAN Configuration.”

- 1 Create a PACKET-FLOWS profile for rate-limiting traffic over the DSL interface of a VLAN circuit. For example:

```
admin> new packet-flows ratelimit-wan
admin> set flow 1 packet-classifier destination-ip-address = 224.0.0.0/4
admin> set flow 1 packet-classifier destination-ip-netmask = 240.0.0.0
admin> set flow 1 scheduling queue-priority = 7
admin> set flow 2 packet-classifier source-ip-address = 10.0.0.0/8
admin> set flow 2 packet-classifier source-ip-netmask = 255.0.0.0
admin> set flow 2 scheduling queue-priority = 6
admin> set flow 2 scheduling yellow-queue-priority = 2
admin> set flow 2 scheduling committed-information-rate = 100
admin> set flow 2 scheduling committed-burst-size = 2760
admin> set flow 2 scheduling excess-burst-size = 27600
admin> set flow 3 scheduling queue-priority = 1
admin> set flow 3 scheduling yellow-queue-priority = 0
admin> set flow 3 scheduling committed-information-rate = 20
admin> set flow 3 scheduling committed-burst-size = 3000
admin> set flow 3 scheduling excess-burst-size = 4500
admin> write -f
```

- 2 Attach the PACKET-FLOWS profile to the CONNECTION profile for the ATM PVC side of the VLAN circuit. For example:

```
admin> read connection user-1
admin> set bridging-options packet-flows = ratelimit-wan
admin> write -f
```

For more complete examples of rate-limiting VLAN traffic, see “Examples of configuring QoS” on page 5-28.

Packet marking

For routed traffic, the IPv4 packet header ToS field can be set to pass priority information on to the next hop. This is useful for IP traffic from a subscriber that is routed to an upstream router.

For bridged and routed VLAN traffic, the 802.1Q VLAN tag can carry both the VLAN ID and a priority value (called a *p-bit value*). The system can set a p-bit value to carry priority information to other p-bit aware devices along the traffic's path.

QoS packet marking for routed traffic

For traffic that is routed through the system, IP ToS marking is supported in both upstream and downstream directions. Values can range from 0 to 255 decimal. The Stinger supports both classification of packets based on this value in incoming packets, and marking of this field in outgoing packets. It does not interpret the value contained in the ToS field, it simply uses it in the configured way. Other network elements, located upstream or downstream, are expected to interpret the values as required, as defined for ToS or DS (for example, as defined in RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*),

Ethernet p-bit marking for bridged or routed VLAN traffic


For traffic that is bridged or routed through the system onto a VLAN in the upstream direction, the Stinger snoops the incoming frames and applies the classification rules to determine the p-bit value to be marked in the VLAN tag.

For routed VLANs, p-bit marking is supported in the upstream direction.

For bridged traffic in the downstream direction, only pretagged stacked VLAN configurations receive VLAN-tagged traffic, so p-bit marking in the downstream direction is restricted to pretagged stacked VLANs.

In Figure 5-2, all CellPipes are operating in LLC-bridged mode. QoS differentiation in the Layer 2 core network is based on priority-bits in the Ethernet header.

Figure 5-2. Sample deployment with VLAN p-bit marking



When transparent bridging is in use, all frames that are flooded are marked with the p-bit value corresponding to the least priority configured (or applicable) for the VLAN interface, and the frames are scheduled on the least priority queue. The same applies to frames with the broadcast MAC destination address. However, if IGMP snooping is not enabled, frames with the multicast MAC destination address are scheduled using configured queues.

If IGMP snooping is enabled, scheduling priority and p-bit marking for multicast frames can be configured at the IP multicast address level in the `packet-classifier` subprofile. Note that a finer granularity of classification (using other fields of the IP header) is not supported for multicast packets.

Overview of packet marking settings

Following are the relevant parameters, shown with default settings:

```
[in PACKET-FLOWS/""]  
marking-type = none  
global-packet-marking = [ { 0 0 0 } { 1 0 0 } { 2 0 0 } { 3 0 0 } { 4 0 0 } {+  
exception-packet-marking = { 8 0 0 }
```

Parameter or subprofile	Purpose
marking-type	Specify whether to mark an IP ToS value (layer 3) or VLAN p-bit value (layer 2) in the packet headers. This setting affects all packet marking for this PACKET-FLOWS profile. Only one type of packet marking (IP ToS or VLAN p-bit) can be configured per profile. With the default <code>none</code> setting, no packet marking is performed. With any other value, you must configure the related field in the <code>global-packet-marking</code> and <code>exception-packet-marking</code> subprofiles. Otherwise, the default marking value of zero is applied. none No packet marking is performed. ip-tos Perform IP ToS marking at the output interface. ethernet-priority Perform Ethernet priority marking at the output interface.
global-packet-marking	An array of 8 indexed subprofiles, one for each scheduling priority in the system. Each of the subprofiles can be used to specify the IP ToS or VLAN p-bit values to be marked in outgoing traffic of the specified priority. Each subprofile contains the following settings: scheduling-priority A read-only value indicating the scheduler priority queue. Packets in this priority queue will be marked with the specified ToS or p-bit value. ip-tos IP ToS value (from 0 to 255) to be set in the ToS field of the IPv4 header. The <code>marking-type</code> parameter must be <code>ip-tos</code> for this value to be marked in outbound packets.

Parameter or subprofile	Purpose
ethernet-priority	802.1P priority value (from 0 to 7) to be set in the VLAN tag of the Ethernet frame. The marking-type parameter must be ethernet-priority for this value to be marked in outbound packets.
exception-packet-marking	A subprofile for specifying the IP ToS or VLAN p-bit values for outgoing packets generated by the Stinger system, such as ARP or ICMP packets.
scheduling-priority	A read-only value indicating the priority queue. Exception packets are always sent on the highest priority queue (8).
ip-tos	IP ToS value (from 0 to 255) to be set in the ToS field of the IPv4 header. The marking-type parameter must be ip-tos for this value to be marked in outbound packets.
ethernet-priority	Currently, for vlan-circuit and vlan-stacked bridged connections, the Stinger does not generate any traffic.

Example of IP ToS marking on a routed VLAN interface

A routed VLAN interface is the router interface or subinterface to which the router's IP address on the VLAN is attached. In a Stinger system, a routed VLAN interface must be mapped to a virtual IP interface. Packets received on a routed VLAN interface are routed based on the IP address, and packets are sent through the routed VLAN interface based on an IP routing decision.

The following steps configure IP ToS marking and priority scheduling on a routed VLAN interface. In this example, the destination IP address is used for packet classification. Each priority class is marked with a specific value in IP ToS field.

- 1 Create a VLAN-ETHERNET profile. This example uses VLAN ID 572.

```
admin> new vlan-ethernet { { 1 8 2 } 572}
admin> set enabled = yes
admin> set bridging-options bridge-type = no-bridging
admin> write -f
```
- 2 Create an IP-INTERFACE profile for the routed VLAN.

```
admin> new ip-interface { { 1 8 2 } 1 }
admin> set ip-address = 10.10.10.10/24
admin> set vlan-enabled = yes
admin> set vlan-id = 572
```

```
admin> set packet-flows = routed-flow
admin> write -f
```

3 Create a PACKET-FLOWS profile for the routed VLAN.

```
admin> new packet-flows routed-flow
admin> set marking-type = ip-tos
admin> set flow 1 packet-classifier destination-ip-address = 20.0.0.0
admin> set flow 1 packet-classifier destination-ip-netmask = 255.0.0.0
admin> set flow 1 scheduling queue-priority = 3
admin> set flow 2 packet-classifier destination-ip-address = 50.0.0.0
admin> set flow 2 packet-classifier destination-ip-netmask = 255.0.0.0
admin> set flow 2 scheduling queue-priority = 2
admin> set flow 3 packet-classifier destination-ip-address = 70.0.0.0
admin> set flow 3 packet-classifier destination-ip-netmask = 255.0.0.0
admin> set flow 3 scheduling queue-priority = 1
admin> set global-packet-marking 2 ip-tos = 1
admin> set global-packet-marking 3 ip-tos = 2
admin> set global-packet-marking 4 ip-tos = 3
admin> write -f
```

Example of Ethernet p-bit marking

The following steps configure Ethernet p-bit marking and priority scheduling on a bridged VLAN output interface. In this example, the destination IP address is used for packet classification. Each priority class is marked with a specific value in IP ToS field.

- 1 Create a VLAN-ETHERNET profile. This example uses VLAN ID 624.

```
admin> new vlan-ethernet { { 1 8 2 } 624 }
admin> set enabled = yes
admin> set bridging-group = 7
admin> set bridge = yes
admin> set bridge-type = transparent-bridging
admin> set bridging-options packet-flows = bridged-flow
admin> write -f
```
- 2 Create a PACKET-FLOWS profile for the bridged VLAN.

```
admin> new packet-flows bridged-flow
admin> set marking-type = ethernet-priority
admin> set flow 1 packet-classifier destination-ip-address = 20.0.0.0
admin> set flow 1 packet-classifier destination-ip-netmask = 255.0.0.0
admin> set flow 1 scheduling queue-priority = 3
admin> set flow 2 packet-classifier destination-ip-address = 50.0.0.0
admin> set flow 2 packet-classifier destination-ip-netmask = 255.0.0.0
admin> set flow 2 scheduling queue-priority = 2
```

```
admin> set flow 3 packet-classifier destination-ip-address = 70.0.0.0
admin> set flow 3 packet-classifier destination-ip-netmask = 255.0.0.0
admin> set flow 3 scheduling queue-priority = 1
admin> set global-packet-marking 2 ethernet-priority = 1
admin> set global-packet-marking 3 ethernet-priority = 2
admin> set global-packet-marking 4 ethernet-priority = 3
admin> write -f
```

Example of mapping ATM QoS to a packet marking value

The following steps show how to use interface grouping to mark packets based on connections' atm-qos service category. The intent is to assign the same interface group number to connections with the same atm-qos service category. For example, CBR connections are placed in interface group 70, VBR connections in interface group 50, and UBR connections in interface group 10. You can then use the interface group to assign marking values. For example, all CBR traffic to be marked with ip-tos = 7, VBR traffic with ip-tos = 5 and UBR with ip-tos = 0.

- 1 Modify CONNECTION profiles to assign interface group numbers based on the connections' service category. For example:

```
admin> read connection cbr-user1
admin> set ip-options qos-interface-group = 70
admin> write -f
admin> read connection vbr-user1
admin> set ip-options qos-interface-group = 50
admin> write -f
admin> read connection ubr-user1
admin> set ip-options qos-interface-group = 10
admin> write -f
```
- 2 Create a PACKET-FLOWS profile that maps the interface groups to appropriate ip-tos marking values.

```
admin> new packet-flows qos-tos-marking
admin> set marking-type = ip-tos
admin> set flow 1 packet-classifier interface-group = 70
admin> set flow 1 scheduling priority = 7
admin> set flow 2 packet-classifier interface-group = 50
admin> set flow 2 scheduling priority = 5
admin> set flow 3 packet-classifier interface-group = 10
admin> set flow 3 scheduling priority = 1
admin> set global-packet-marking 8 ip-tos = 7
admin> set global-packet-marking 6 ip-tos = 5
admin> set global-packet-marking 2 ip-tos = 0
admin> write -f
```

QoS-related connection and interface settings

The following parameters, shown with default values, are related to QoS handling. Note that the VLAN-ETHERNET scheduling subprofile applies only to bridged VLAN interfaces. For routed VLAN interfaces, the scheduling subprofile in the IP-INTERFACE profile must be used.

```
[in CONNECTION/"" :ip-options]
qos-interface-group = 0
packet-flows = ""

[in CONNECTION/"" :bridging-options]
qos-interface-group = 0
packet-flows = ""

[in CONNECTION/"" :atm-qos-options]
usr-up-stream-contract = default
usr-dn-stream-contract = default

[in VLAN-ETHERNET/{ { shelf-1 first-control-module 2 } n } :bridging-options]
qos-interface-group = 0
packet-flows = ""

[in IP-INTERFACE/{ { shelf-1 first-control-module 2 } 0 } ]
qos-interface-group = 0
packet-flows = ""

[in VLAN-ETHERNET/{ { shelf-1 first-control-module 2 } n } :scheduling]
priority = 0
bandwidth = 0

[in IP-INTERFACE/{ { shelf-1 first-control-module 2 } 0 } :scheduling]
priority = 0
bandwidth = 0
```

Parameter

qos-interface-group

Setting

Numeric field, from 0 to 255, used to group input interfaces for QoS handling. The assigned value can be used as a classifier. For an example configuration using interface grouping, see “Prioritizing IP packet flows based on DSL service contracts” on page 5-29.

For routed subscriber interfaces, set the group number in the CONNECTION ip-options subprofile. For bridged subscriber interfaces, set the number in the connection bridging-options subprofile. For routed VLANs, the qos-interface-group is set in the IP-INTERFACE profile and in case of bridged VLANs, it is set in the VLAN-ETHERNET profile.

Parameter	Setting
packet-flows	Name of the PACKET-FLOWS profile to be attached to the output interface. For routed subscriber output interfaces, attach the PACKET-FLOWS profile in the CONNECTION ip-options subprofile. For bridged subscriber output interfaces, attach the profile in the CONNECTION bridging-options subprofile. For routed VLANs, the PACKET-FLOWS profile can be set in the IP-INTERFACE profile. For bridged VLANs, it is set in the VLAN-ETHERNET profile. To attach a PACKET-FLOWS profile to a RADIUS-authenticated interface, use the Ascend-Packet-Flows (324) vendor-specific attribute.
usr-up-stream-contract	Names of the ATM-QOS profiles to be applied to upstream traffic and downstream traffic, respectively. See “ATM QoS and IP QoS considerations” on page 5-24.
usr-dn-stream-contract	
priority	Scheduler priority of the virtual IP or VLAN Ethernet interface, from 0 to 3. The default value of 0 indicates the lowest priority.
bandwidth	Scheduling rate of the virtual IP or VLAN Ethernet interface, expressed in Kbps, from 0 to 1,000,000. The default value of 0 indicates full bandwidth. To configure a nonzero number, the priority setting must also be nonzero. For IP2000 versions A and B with a revision level of 25 or less, the minimum scheduling rate is 801Kbps. For IP2000 version B with a revision level of 26 and higher, the minimum scheduling rate is 401Kbps. The system ensures that these minimums are respected.

Applying a PACKET-FLOWS profile to an output interface

The following commands apply a configured PACKET-FLOWS profile named `demo-priority-flow` on a Gigabit Ethernet interface:

```
admin> read ip-interface { { 1 8 2 } 0 }  
admin> set packet-flows = demo-priority-flow  
admin> write -f
```

Inheritance of PACKET-FLOWS configurations on virtual IP interfaces

Both WAN and LAN interfaces can have virtual IP interfaces, although only LAN interfaces support static configuration of a virtual interface.

Dynamic virtual interfaces, such as those created on a WAN interface when relaying DHCP requests, always inherit the same `packet-flow` configuration and interface group as the main IP interface on the port.

Packet flow processing occurs at the output side of an interface, and is applicable to all statically configured virtual interfaces. A virtual interface and the main interface on the same port do not have to use the same PACKET-FLOWS profile.

Virtual IP interfaces and interface grouping

When a port has been assigned a `qos-interface-group` number, the system must be able to determine which virtual IP interface on the port is receiving a packet flow, because interface grouping occurs at the input side of the interface. Because it is not ensured that the system can determine the correct virtual interface carrying a packet flow, a virtual IP interface must use the same interface group as the main IP interface on the port. If a different interface group is specified in the virtual interface IP-INTERFACE profile, the system logs a warning message applies the group number of the main IP interface.

ATM QoS and IP QoS considerations

The ATM-QOS and PACKET-FLOWS profiles can be used in conjunction to configure connection bandwidth and the per-flow bandwidth. The ATM-QOS profile definitions are used in the downstream direction to specify the ATM QoS contract (CBR, nrtVBR, rtVBR, or UBR), which sets the priority level of the connection, with CBR at the highest priority and UBR at the lowest. This connection-level prioritization is particularly useful when several connections share the same physical link. The ATM-QOS profile traffic parameters (PCR for CBR connections, and SCR for VBR connections) are used to shape the downstream traffic at the specified rates. In the upstream direction, ATM policing performed on the traffic received on the connection.

Using an unlimited rate contract for an ATM connection can have an impact regarding priorities defined by the PACKET-FLOWS profile. For example, suppose a CONNECTION profile uses the default ATM QoS, which applies a UBR contract, and applies a PACKET-FLOWS profile that defines three flows mapped to three different priorities, with flow 1 at the highest priority and a throughput of 2Mbps, flow 2 at medium priority and a throughput of 4Mbps, and flow 3 at the lowest priority and a throughput of 5Mbps. Suppose the connection has a line rate of 6Mbps. Because there is no traffic shaper associated with the connection, the network processor sends the full 11Mbps to the DSL line, causing discards at the LIM side regardless of the traffic priority.

Configuring Ethernet IP traffic shaping

The `scheduling` subprofile of the IP-INTERFACE profile enables you to apply a traffic shaper and a priority for a logical interface on a Gigabit Ethernet port. The `scheduling` configuration is applicable to all kinds of routed connections, including routed VLANs.



Note For IP2000 version B with a revision level of 26 and higher, the minimum scheduling rate is 401Kbps. For IP2000 versions A and B with a revision level of 25 or less, the minimum scheduling rate is 801Kbps.

The following commands specify middle priority (rather than the default low priority) for traffic on the IP interface of a Gigabit Ethernet port:

```
admin> read ip-interface { { 1 8 2 } 0 }  
admin> set scheduling priority = 2
```



```
admin> write -f
```

Configuring VLAN Ethernet traffic shaping

VLAN traffic shaping is typically used when different kinds of traffic on a network (such as data, VoIP, and video) are each bridged to different VLAN Ethernet interfaces, and the service provider needs control over the amount of bandwidth each type of traffic can use. VLAN traffic shaping is supported only for VLAN Ethernet interfaces. VLAN traffic shaping is supported as shown in Table 5-6:

Table 5-6. VLAN traffic shaping configurations

Type of VLAN	Traffic shaping description
1:1 VLAN bridging	Traffic shaping is performed individually for each VLAN. Since each VLAN corresponds to one ATM PVC, traffic from each user can be shaped independently at the <code>vlan-ethernet</code> interface.
N:1 VLAN bridging (transparent bridged interfaces)	Traffic shaping is performed individually for each VLAN interface. The aggregate traffic from multiple user connections can be shaped as specified by the parameters configured for the VLAN interface.
Stacked VLANs	Traffic shaping is performed at the NSP VLAN level. All traffic transmitted over an NSP VLAN shares the same set of class-based queues. The aggregate traffic from multiple user connections can be shaped as specified by the parameters configured for the NSP VLAN. This is applicable to both untagged and pre-tagged stacked-VLAN connections.

The following commands specify a scheduler of priority 1 and traffic shaping that transmits PDU traffic at 500 kbps for traffic on the VLAN Ethernet interface:

```
admin> new vlan-ethernet { { 1 8 2 } 500 }
admin> set enabled = yes
admin> set bridging-options bridging-group = 7
admin> set bridging-options bridge = yes
admin> set bridging-options bridge-type = transparent-bridging
admin> set scheduling priority = 1
admin> set scheduling bandwidth = 500
admin> write -f
```

QoS-related settings in the SYSTEM profile

Settings in the SYSTEM profile affect call management systemwide, and are documented in the *Stinger Administration Guide*. The next sections provide overviews of settings that can affect performance when QoS is applied.

Performance recommendations

When PACKET-FLOWS or FILTER profiles have been applied to output interfaces, the Stinger system software must communicate resulting changes in traffic processing requirements to the underlying system hardware. These updates to the hardware are

typically time intensive, so creating a batch of these requests is recommended to improve the overall performance of the system (system performance) and to control the rate at which connections or service is restored to end-users (call performance).

The following parameters, shown with default values, provide performance enhancements for interactions with the network processor:

```
[in SYSTEM]
np-update-time = 0
np-default-filtering-policy = forward
np-fpp-compact-timer = 3600
```

Setting `np-update-time` to between 10-20 seconds is strongly recommended, particularly when a large number of filter and/or classification rules are configured on the system. For details about `np-update-time` and the other network processor performance options, see the *Stinger Administration Guide*.

Configurable queue size for IPTV traffic on DSL links

For a system heavily loaded with subscriber uplinks to a Gigabit Ethernet interface, the system can begin to discard traffic sent to the LIM based on scheduler memory and queue size. For DSL links, the default queue is 200 blocks (with 256 bytes per block). For trunk connections and Gigabit Ethernet links, the queue size is not configurable.

You can increase the queue size for all LIM interfaces system wide by setting the following parameter, shown with its default value:

```
[in SYSTEM]
scheduler-queue-size = 200
```

For IPTV applications, you should increase the default `scheduler-queue-size` to 500, as shown in the following commands:

```
admin> read system
admin> set scheduler-queue-size = 500
admin> write -f
```



Note When PACKET-FLOWS profiles are in use, the maximum setting for the `scheduler-queue-size` parameter in the SYSTEM profile is 1000. When the `scheduler-queue-size` parameter is set to a value greater than 1000, the system does not conform to the priority specified in the PACKET-FLOWS profile.

Tracking rate adaptation for downstream traffic shaping

The IP CoS subsystem calculates new QoS values for a DSL line if the rate drops below the configured peak cell rate (PCR) or sustainable cell rate (SCR) for a connection. As soon as the line rate adapts back to a higher rate, the original QoS values are reapplied. In this way, traffic shaping can consider assigned CoS priorities while discarding traffic at the IP layer itself, before the packets are fragmented into cells.

This optimization is intended for use only for a configuration using one underlying ATM circuit per DSL line (for PPPoA or PPPoE sessions) or one BIR connection per DSL line. It applies only to non-UBR traffic on high-bandwidth (HB) Annex A LIMs with ADSL2+ capability (STGR-LIM-AP-72-HB or STGR-LIM-AP-48-HB) or VDSL LIMs (STGR-LIM-VD-24). It takes effect only when the training rate of a DSL line

goes below the configured PCR or SCR and is decreasing. It does not take effect if the training rate drops to zero.

How traffic shaping typically occurs

The control module (CM) network processor (NP) shapes downstream traffic by discarding packets received at a higher rate than a connection's expected ATM QoS. The CM NP selects which packets to discard based on IP CoS priorities as configured in a PACKET-FLOWS profile.

After discarding packets that exceed the expected ATM QoS, the system ships the downstream traffic to a LIM as a cell stream. Because the CM NP has already discarded packets based on priority, the LIM does not have to discard excess cells.

How rate adaptation affects traffic shaping

ADSL2+ uses seamless rate adaptation (SRA) to adapt the downstream modulation based on ambient conditions, to ensure that a sufficient signal-to-noise ratio (SNR) is maintained and the line does not drop if the ambient conditions get bad. This type of rate adaptation can happen at any time, and the rate to which the ADSL2+ line adapts can be lower than the expected QoS. However, the CM NP does not track the line rate transition, so it continues to shape traffic based on the configured QoS. The LIM then discards the excess cells to the lower line rate. Naturally, cell discards are done without regard to the priority of the traffic.

How to enable the rate-tracking optimization

The following parameter, shown with its default value, causes the system to track the underlying DSL line speed of a connection. If the DSL line speed adapts to a rate lower than the PCR or SCR, the system dynamically realigns the QoS of the connection with the underlying DSL line speed. As soon as the line adapts back to its higher rate, the configured QoS is reinstated.

```
[in ATM-QOS/""]  
track-underlying-dsl = no
```

The following commands enable the feature:

```
admin> read atm-qos dn-q  
admin> set track-underlying-dsl = yes  
admin> write -f
```

How this feature works on the VDSL LIM

On the VDSL LIM, all traffic is shaped by the CM NP only. Whether it is processing unicast traffic, multicast traffic, or a combination of the two, it discards packets on the basis of the PACKET-FLOWS priorities. CoS optimization offers an advantage on the VDSL LIM for any kind of traffic.

How this feature works on the ADSL2+ HB LIM

On the ADSL2+ HB LIM, CoS optimization offers an advantage for pure unicast traffic. When distributed multicast is in use, this optimization does not degrade quality of service for pure multicast traffic or multicast with unicast traffic.

Traffic processing changes with distributed multicast

With distributed multicast, the CM sends the ADSL2+ HB LIM a single stream of multicast traffic and the LIM duplicates this stream over each DSL line in the multicast group. So, the LIM actually creates the multicast IP packets. When the CM NP shapes unicast traffic according to the new line rate, it hasn't yet accounted for the multicast streams being replicated on the LIM. For this reason, unicast traffic suffers a performance penalty if it has a lower priority than multicast traffic and it is being transmitted on a multicast-enabled ADSL2+ HB LIM. For information on unicast and multicast priorities with distributed multicast on an ADSL2+ HB LIM, see the section on distributed multicast in the *Stinger IP Control Module Configuration Guide*.

Effect on how ADSL2+ HB LIMs prioritize packets

When only unicast traffic is being transmitted, configured CoS priorities determine which traffic is discarded. Each and every priority configured in the PACKET-FLOWS profile is honored.

With pure multicast traffic or a combination of multicast and unicast traffic, the available bandwidth is used to send all of the highest priority traffic "type," which is determined as being either *unicast* or *multicast*. Once the highest priority type is determined, no finer distinction is made between two streams of the same type.

If no bandwidth remains thereafter, the other traffic type (the type that is *not* highest priority) will be discarded. If bandwidth still remains, it be used to send the other traffic type.

The system determines which type of traffic will be determined as highest priority as follows:

- If no priority is configured, multicast is given highest priority.
- If a unicast and multicast stream both share the same highest priority number, then multicast is given highest priority.
- If a unicast stream has the highest priority, then all unicast streams are given priority over multicast.
- If a multicast stream has the highest priority, then all multicast streams are given priority over unicast.

Examples of configuring QoS


This section contains the following configuration examples:

- Prioritizing traffic based on the level of end-user service contract
- Classifying several types of IP traffic for prioritization on one ATM PVC
- Using packet flow prioritization together with ATM QoS to prioritize bandwidth use for different types of IP traffic
- Rate limiting and traffic shaping with N:1 VLAN bridging
- Rate limiting and traffic shaping with VLAN stacking
- Rate limiting and traffic shaping for 1:1 VLAN bridging
- Sample PPPoE connection with and without line rate tracking
- Sample BIR connection with line rate tracking

Prioritizing IP packet flows based on DSL service contracts

In this example, the provider configures processing priorities according to the price of the end-user DSL contract. For a low-price service contract (bronze), packet flows are assigned a low processing priority, for a medium-price contract (silver), packet flows will be processed at medium priority, and for a high-price contract (gold), packet flows will have high-priority processing. The PACKET-FLOWS profile is applied to the Ethernet output interface.

Figure 5-3. Using interface grouping to prioritize traffic by service level



The result of this configuration is that the traffic coming from a user with a gold service contract is output on the Ethernet interface from the high-priority queue. Once the high-priority queue is empty, traffic from a user with a silver service contract is output from the medium-priority queue. When the high-priority and medium-priority queues are empty, traffic from a user with a bronze contract is transmitted on the output interface.

This sample configuration requires the following steps:

1. Configure subscriber profiles with the appropriate qos-interface-group value. In this example, bronze customers are assigned to group 1, silver customers are assigned to group 2 and gold customers are assigned to interface group 3.

```
admin> read connection user-gold-1
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 192.168.100.1/32
admin> set ip-options local-address = 192.168.100.200/32
admin> set ip-options qos-interface-group = 3
admin> set atm-options nailed-group = 21
admin> write -f
admin> new connection user-silver-1
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 192.168.100.10/32
admin> set ip-options local-address = 192.168.100.201/32
admin> set ip-options qos-interface-group = 2
admin> set atm-options nailed-group = 2
admin> write -f
```


- ```
admin> new connection user-bronze-1
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 192.168.100.20/32
admin> set ip-options local-address = 192.168.100.202/32
admin> set ip-options qos-interface-group = 1
admin> set atm-options nailed-group = 11
admin> write -f
```
- 2 Configure the PACKET-FLOWS profile to classify and prioritize traffic by interface group. Traffic from bronze customers will have priority 1, from silver customers will have priority 3, and from gold customers will have priority 5.
- ```
admin> new packet-flows service-contracts
admin> set flow 1 packet-classifier interface-group = 1
admin> set flow 1 scheduling queue-priority = 1
admin> set flow 2 packet-classifier interface-group = 2
admin> set flow 2 scheduling queue-priority = 3
admin> set flow 3 packet-classifier interface-group = 3
admin> set flow 3 scheduling queue-priority = 5
admin> write -f
```
- 3 Apply the PACKET-FLOWS profile to the output interface (Gigabit Ethernet).
- ```
admin> read ip-interface { { 1 8 2 } 0 }
admin> set ip-address = 10.60.1.1/8
admin> set netmask = 255.0.0.0
admin> set packet-flows = service-contracts
admin> write -f
```

### ***Prioritizing different kinds of IP traffic on an ATM PVC***

In this example, the provider needs to classify different kinds of IP traffic across a single ATM PVC. The PACKET-FLOWS profile will be applied to the DSL connection as the output interface.

For prioritization changes related to distributed multicast with high-bandwidth LIMs, see “Changes in packet-flow prioritization with distributed multicast” on page 9-29.

Figure 5-4. Unicast and multicast video share the same priority



The provider classifies IP traffic types as shown in Table 5-7. All control traffic such as IGMP, RIP, and so forth, is assigned the highest priority and has precedence over the scheduled QoS priorities.

Table 5-7. Sample IP traffic types and priorities across an ATM PVC

| Kind of IP traffic               | Priority            | Classification                     |
|----------------------------------|---------------------|------------------------------------|
| Multicast video                  | Medium (6)<br>video | Multicast destination address      |
| Unicast video pay per view (PPV) | Medium (6)          | PPV server source IP address       |
| Data traffic                     | Low (5)             | Default rule for all other traffic |



**Note** Because multicast video and the unicast PPV have the same queue-priority value, they will share the same queue. See “Changes in packet-flow prioritization with distributed multicast” on page 9-29 for related information.

This sample configuration requires the following basic steps:

1. Configure the PACKET-FLOWS profile to classify and prioritize traffic by type. Multicast video and PPV are assigned medium priority, and data traffic is assigned low priority (by defining a default rule). PPV traffic is classified by the source IP address of the PPV server (192.168.200.200).

```
admin> new packet-flows
admin> set name = mixed-traffic
admin> set flow 1 packet-classifier destination-ip-address = 224.0.0.0
admin> set flow 1 packet-classifier destination-ip-netmask = 240.0.0.0
admin> set flow 1 scheduling queue-priority = 6
admin> set flow 2 packet-classifier source-ip-address = 192.168.200.200
admin> set flow 2 packet-classifier source-ip-netmask = 255.255.255.255
admin> set flow 2 scheduling queue-priority = 6
admin> set flow 3 scheduling queue-priority = 5
admin> write -f
```


- 2 Apply the PACKET-FLOWS profile to the output interface (the DSL side).

```
admin> new connection user-1
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 192.168.100.40/32
admin> set ip-options local-address = 192.168.100.204/32
admin> set ip-options packet-flows = mixed-traffic
admin> set atm-options nailed-group = 4
admin> write -f
```

## Prioritizing traffic using both IP and ATM QoS

In this example, the provider needs to classify different kind of IP traffic across two ATM PVCs on the same DSL interface. The PACKET-FLOWS profile is applied to the DSL connection as the output interface, and ATM QoS is applied to each ATM PVC separately.

Figure 5-5. Prioritizing traffic at the connection level and flow level



In this example, two PACKET-FLOWS profiles are configured, because the two ATM PVCs have very different classification requirements, as shown in Table 5-8:

Table 5-8. Sample IP traffic types and priorities across two PVCs

| ATM PVC   | Kind of IP traffic | Priority   | Classification                     |
|-----------|--------------------|------------|------------------------------------|
| video-pvc | Signaling          | High (7)   | All other traffic                  |
|           | Multicast video    | Medium (6) | Multicast destination address      |
|           | Unicast video/PPV  | Medium (6) | PPV server source IP address       |
|           | Miniguide          | Low (5)    | TCP protocol (for HTTP)            |
| data-pvc  | Web TV             | High (7)   | UDP protocol (for RTP)             |
|           | Internet traffic   | Low (6)    | Default rule for all other traffic |



**Caution** The setup shown in Table 5-8 is just an example. Because it grants priority 7 to a classification for “all other traffic” (the default rule), it could cause the system to grant high priority to unauthorized traffic such as packets received during a denial-



of-service attack or a Blaster worm attack. This type of security issue must be considered before configuring a packet flow with a high-priority default rule.

This sample configuration requires the following basic steps:

- 1 Configure a PACKET-FLOWS profile for **video-pvc**. The classification for packet flow 1 defines multicast video at medium priority. Packet flow 2 specifies unicast PPV at medium priority, on the basis of the source IP address of the PPV server (192.168.10.1).

Packet flow 3 classifies HTTP traffic on the basis of the TCP protocol. A null classification (a default rule) for flow 4, specifies that all other IP traffic (such as signaling traffic) carried on the PVC should be assigned the highest priority.

```
admin> new packet-flows
admin> set name = video-traffic
admin> set flow 1 packet-classifier destination-ip-address = 224.0.0.0/8
admin> set flow 1 packet-classifier destination-ip-netmask = 255.0.0.0
admin> set flow 1 scheduling queue-priority = 6
admin> set flow 2 packet-classifier source-ip-address = 192.168.10.1/32
admin> set flow 2 packet-classifier source-ip-netmask = 255.255.255.255
admin> set flow 2 scheduling queue-priority = 6
admin> set flow 3 packet-classifier ip-protocol = 6
admin> set flow 3 scheduling queue-priority = 5
admin> set flow 4 scheduling queue-priority = 7
admin> write -f
```

- 2 Configure a PACKET-FLOWS profile for **data-pvc**, by classifying packet flow 1 for WebTV at high priority, and then defining a default rule for flow 2 with a priority of 6 for all other IP traffic carried on the PVC:

```
admin> new packet-flows
admin> set name = internet-traffic
admin> set flow 1 packet-classifier ip-protocol = 17
admin> set flow 1 scheduling queue-priority = 7
admin> set flow 2 scheduling queue-priority = 6
admin> write -f
```

- 3 Configure three ATM-QOS profiles, one for downstream traffic for each PVC, and one to be used by both PVCs for upstream traffic.

In this example, the DSL line supporting the two PVCs provides 6Mbps downstream and 256Kbps upstream bandwidth. The provider has to limit the bandwidth accessible to the two interfaces, as shown in Table 5-9:

*Table 5-9. Sample bandwidth limitations*

| ATM PVC   | Bandwidth          | Service classification |
|-----------|--------------------|------------------------|
| video-pvc | 5.5Mbps downstream | rtVBR at 5500Kbps      |
|           | 128k upstream      | nrtVBR at 128Kbps      |

Table 5-9. Sample bandwidth limitations (Continued)

| ATM PVC  | Bandwidth          | Service classification |
|----------|--------------------|------------------------|
| data-pvc | 320Kbps downstream | rtVBR at 320Kbps       |
|          | 128k upstream      | nrtVBR at 128Kbps      |

The next set of commands defines an ATM-QOS profile for data-pvc downstream traffic:

```
admin> new atm-qos
admin> set contract-name = internet-downstream
admin> set traffic-descriptor-type = noclpscr-cdvt
admin> set atm-service-category = real-time-vbr
admin> set peak-rate-kbits-per-sec = 320
admin> set peak-cell-rate-cells-per-sec = 754
admin> set sustainable-rate-kbits-per-sec = 320
admin> set sustainable-cell-rate-cells-per-sec = 754
admin> set ignore-max-burst-size = no
admin> set max-burst-size = 1
admin> set aal-type = aal-5
admin> set early-packet-discard = yes
admin> write -f
```

The next set of commands defines an ATM-QOS profile for video-pvc downstream traffic:

```
admin> new atm-qos
admin> set contract-name = video-downstream
admin> set traffic-descriptor-type = noclpscr-cdvt
admin> set atm-service-category = real-time-vbr
admin> set peak-rate-kbits-per-sec = 5500
admin> set peak-cell-rate-cells-per-sec = 12971
admin> set sustainable-rate-kbits-per-sec = 5500
admin> set sustainable-cell-rate-cells-per-sec = 12971
admin> set ignore-max-burst-size = no
admin> set max-burst-size = 1
admin> set aal-type = aal-5
admin> set early-packet-discard = yes
admin> write -f
```

Because both PVCs support the same amount of upstream bandwidth, they can both use the same ATM-QOS profile. The next set of commands defines an ATM-QOS profile for video-pvc and data-pvc upstream traffic:

```
admin> new atm-qos
admin> set contract-name = upstream
```

```
admin> set traffic-descriptor-type = noclpscr-cdvt
admin> set atm-service-category = real-time-vbr
admin> set peak-rate-kbits-per-sec = 128
admin> set peak-cell-rate-cells-per-sec = 301
admin> set sustainable-rate-kbits-per-sec = 128
admin> set sustainable-cell-rate-cells-per-sec = 301
admin> set ignore-max-burst-size = no
admin> set max-burst-size = 1
admin> set aal-type = aal-5
admin> set early-packet-discard = yes
admin> write -f
```

- 4 Apply the PACKET-FLOWS and ATM-QOS profiles to the output interfaces (the subscriber ATM PVCs in this example).

The following commands configure video-pvc:

```
admin> new connection
admin> set station = user-1-video-pvc
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 192.168.100.100/32
admin> set ip-options local-address = 192.168.100.210/32
admin> set ip-options packet-flows = video-traffic
admin> set atm-options nailed-group = 6
admin> set atm-qos-options usr-up-stream-contract = upstream
admin> set atm-qos-options usr-dn-stream-contract = video-downstream
admin> write -f
```


The next commands configure data-pvc:

```
admin> new connection
admin> set station = user-1-data-pvc
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 192.168.100.101/32
admin> set ip-options local-address = 192.168.100.211/32
admin> set ip-options packet-flows = internet-traffic
admin> set atm-options vci = 36
admin> set atm-options nailed-group = 6
admin> set atm-qos-options usr-up-stream-contract = upstream
admin> set atm-qos-options usr-dn-stream-contract = internet-downstream
admin> write -f
```

## Configuring bridging VLAN Ethernet QoS

In this example, the Stinger IP DSLAM receives pre-tagged VLAN traffic from the CellPipe®. The incoming VLAN traffic is p-bit marked, and the system must mark the p-bit again before sending the stacked frames out to NSP 1. This sample setup is shown in Figure 5-6.

Figure 5-6. Stacked VLAN requiring p-bit remarking on the output interface



The incoming p-bit values from the CellPipe® and the p-bit values in the output frames to NSP VLAN 1 are shown in Table 5-10.

Table 5-10. Ethernet p-bit remarking table

| p-bit value in input frame | p-bit value in output frame |
|----------------------------|-----------------------------|
| 1                          | 1                           |
| 2                          | 3                           |
| 3                          | 5                           |

To configure the p-bit priority mappings, follow these steps:


- 1 Make sure bridging is enabled in the ETHERNET profile. (If the layer-2 core network is using a specific EtherType value for stacked VLAN frames, you must also set the vlan-stack-tag-type parameter to that value.)  
admin> read ethernet { 1 8 2 }  
admin> set bridging-enabled = yes  
admin> write -f
- 2 Create a PACKET-FLOWS profile that classifies the traffic by incoming p-bit value and reassigns a new p-bit value, as shown in Table 5-10.  
admin> new packet-flows  
admin> set name = pbitremark  
admin> set marking-type = ethernet-priority  
admin> set flow 1 layer2-classifier ethernet-priority = 1  
admin> set flow 1 scheduling queue-priority = 1  
admin> set flow 2 layer2-classifier ethernet-priority = 2  
admin> set flow 2 scheduling queue-priority = 2  
admin> set flow 3 layer2-classifier ethernet-priority = 3

- ```
admin> set flow 3 scheduling queue-priority = 3
admin> set global-packet-marking 2 ethernet-priority = 1
admin> set global-packet-marking 3 ethernet-priority = 3
admin> set global-packet-marking 4 ethernet-priority = 5
admin> write -f
```
- 3 Create the NSP VLAN and apply the PACKET-FLOWS profile. In this example, the VLAN is assigned VLAN ID 572 and bridging-group 300.
- ```
admin> new vlan-ethernet
admin> set interface-address physical shelf = shelf-1
admin> set interface-address physical slot = first-control-module
admin> set interface-address physical item-number = 2
admin> set interface-address logical-item = 572
admin> set enabled = yes
admin> set bridging-options bridging-group = 300
admin> set bridging-options bridge = yes
admin> set bridging-options bridge-type = stacked-vlan
admin> set bridging-options packet-flows = pbitremark
admin> write -f
```
- 4 Configure the flow-services mapping between the incoming (enterprise) VLAN tags on the DSL interface and user ID values that are unique within the NSP VLAN. Then create a CONNECTION profile for the inbound CellPipe® connection with the same bridging-group number as the NSP VLAN, and apply the FLOW-SERVICES profile. For details about configuring stacked VLAN, see “Configuring stacked VLANs” on page 3-27.

## Configuring traffic shaping and rate limiting for a N:1 VLAN bridging

Figure 5-7 shows a Stinger IP DSLAM transparent-bridging multiple user PVCs onto a VLAN. In this configuration, PDU traffic will be transmitted across the VLAN Ethernet interface at 500Kbps with priority 1. The DSL interfaces will transmit multicast traffic at an unlimited rate, but will rate limit unicast and all other traffic flows.

Figure 5-7. Sample N:1 VLAN bridging with traffic shaping and rate limiting



To configure this N:1 VLAN bridging setup to rate-limit traffic transmitted over DSL interfaces, and to shape traffic over the VLAN Ethernet interface, follow these steps:

- 1 If bridging is not enabled on the Ethernet port, enable it as described in “Enabling layer 2 bridging for VLAN operations” on page 2-3.

- 2 Create a bridge group for VLAN 50. For example:

```
admin> new bridge-group
admin> set enable = yes
admin> set bridging-group = 7
admin> set lan-router-interface physical shelf = shelf-1
admin> set lan-router-interface physical slot = first-control-module
admin> set lan-router-interface physical item = 2
admin> set lan-router-interface logical-item = 50
admin> write -f
```
- 3 Configure a VLAN-ETHERNET profile for VLAN 50 with a scheduler of priority 1 and traffic shaping that transmits PDU traffic at 500 kbps.

```
admin> new vlan-ethernet
admin> set interface-address physical shelf = shelf-1
admin> set interface-address physical slot = first-control-module
admin> set interface-address physical item-number = 2
admin> set interface-address logical-item = 50
admin> set enabled = yes
admin> set bridging-options bridging-group = 7
admin> set bridging-options bridge = yes
admin> set bridging-options bridge-type = transparent-bridging
admin> set scheduling priority = 1
admin> set scheduling bandwidth = 500
admin> write -f
```
- 4 To configure rate limiting for DSL interfaces, create a PACKET-FLOWS profile with an unlimited multicast flow, followed by a rate-limited unicast flow and a rate-limited default flow. For example:

```
admin> new packet-flows
admin> set name = ratelimit-wan
admin> set flow 1 packet-classifier destination-ip-address = 224.0.0.0/4
admin> set flow 1 packet-classifier destination-ip-netmask = 240.0.0.0
admin> set flow 1 scheduling queue-priority = 7
admin> set flow 2 packet-classifier source-ip-address = 10.0.0.0/8
admin> set flow 2 packet-classifier source-ip-netmask = 255.0.0.0
admin> set flow 2 scheduling queue-priority = 6
admin> set flow 2 scheduling yellow-queue-priority = 2
admin> set flow 2 scheduling committed-information-rate = 100
admin> set flow 2 scheduling committed-burst-size = 3000
admin> set flow 2 scheduling excess-burst-size = 4500
admin> set flow 3 scheduling queue-priority = 1
admin> set flow 3 scheduling yellow-queue-priority = 0
```

```
admin> set flow 3 scheduling committed-information-rate = 20
admin> set flow 3 scheduling committed-burst-size = 3000
admin> set flow 3 scheduling excess-burst-size = 4500
admin> write -f
```

5 Apply the PACKET-FLOWS profile to transparent-bridged CONNECTION profiles. For example:


```
admin> new connection cpe-1
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options ip-routing-enabled = no
admin> set bridging-options bridging-group = 7
admin> set bridging-options bridge = yes
admin> set bridging-options bridge-type = transparent-bridging
admin> set bridging-options packet-flows = ratelimit-wan
admin> set atm-options nailed-group = 251
admin> write -f
```

## Configuring rate limiting and traffic shaping with VLAN stacking

In the sample stacked VLAN setup shown in Figure 5-8, a layer-2 core is partitioned among three NSPs. With the following configuration, the aggregate traffic transmitted over NSP VLAN 478 interface will be rate-limited based on flow classification, and traffic that is passed by the policer will be shaped at 500kbps.

Traffic transmitted in the downstream direction over DSL interfaces is individually rate-limited as specified in the flow classification. That is, traffic towards user-1 and user-2 (for example) are rate limited independently.

Figure 5-8. Stacked VLAN with rate limiting and traffic shaping



To configure the sample rate limiting and traffic shaping with VLAN stacking, follow these steps:

- 1 Create a PACKET-FLOWS profile for rate-limiting traffic over the NSP VLAN interface. For example:  
admin> new packet-flows ratelimit-vlan  
admin> set flow 1 packet-classifier destination-ip-address = 224.0.0.0/4  
admin> set flow 1 packet-classifier destination-ip-netmask = 240.0.0.0  
admin> set flow 1 scheduling queue-priority = 7  
admin> set flow 2 packet-classifier source-ip-address = 10.0.0.0/8  
admin> set flow 2 packet-classifier source-ip-netmask = 255.0.0.0  
admin> set flow 2 scheduling queue-priority = 6  
admin> set flow 2 scheduling yellow-queue-priority = 2  
admin> set flow 2 scheduling committed-information-rate = 100  
admin> set flow 2 scheduling committed-burst-size = 2760  
admin> set flow 2 scheduling excess-burst-size = 27600  
admin> set flow 3 scheduling queue-priority = 1  
admin> set flow 3 scheduling yellow-queue-priority = 0  
admin> set flow 3 scheduling committed-information-rate = 20  
admin> set flow 3 scheduling committed-burst-size = 3000  
admin> set flow 3 scheduling excess-burst-size = 4500  
admin> write -f
- 2 Configure a VLAN-ETHERNET profile for VLAN 478 with a scheduler of priority 1 and traffic shaping that transmits PDU traffic at 500 kbps, and attach the PACKET-FLOWS profile. For example:  
admin> new vlan-ethernet  
admin> set interface-address physical shelf = shelf-1  
admin> set interface-address physical slot = first-control-module  
admin> set interface-address physical item-number = 2  
admin> set interface-address logical-item = 478  
admin> set enabled = yes  
admin> set bridging-options bridging-group = 7  
admin> set bridging-options bridge = yes  
admin> set bridging-options bridge-type = stacked-vlan  
admin> set bridging-options packet-flows = ratelimit-vlan  
admin> set scheduling priority = 1  
admin> set scheduling bandwidth = 500  
admin> write -f
- 3 Create a PACKET-FLOWS profile for rate-limiting traffic over the DSL interfaces. For example:  
admin> new packet-flows



```
admin> set name = ratelimit-wan
admin> set flow 1 packet-classifier destination-ip-address = 224.0.0.0/4
admin> set flow 1 packet-classifier destination-ip-netmask = 240.0.0.0
admin> set flow 1 scheduling queue-priority = 7
admin> set flow 2 packet-classifier source-ip-address = 10.0.0.0/8
admin> set flow 2 packet-classifier source-ip-netmask = 255.0.0.0
admin> set flow 2 scheduling queue-priority = 6
admin> set flow 2 scheduling yellow-queue-priority = 2
admin> set flow 2 scheduling committed-information-rate = 100
admin> set flow 2 scheduling committed-burst-size = 2760
admin> set flow 2 scheduling excess-burst-size = 27600
admin> set flow 3 scheduling queue-priority = 1
admin> set flow 3 scheduling yellow-queue-priority = 0
admin> set flow 3 scheduling committed-information-rate = 20
admin> set flow 3 scheduling committed-burst-size = 3000
admin> set flow 3 scheduling excess-burst-size = 4500
admin> write -f
```

- 4 Create CONNECTION profiles for user-1 and user-2 and attach the PACKET-FLOWS profile:


```
admin> new connection user-1
admin> set encapsulation-protocol = atm
admin> set ip-options ip-routing-enabled = no
admin> set bridging-options bridging-group = 7
admin> set bridging-options bridge = yes
admin> set bridging-options bridge-type = stacked-vlan
admin> set bridging-options vlan-stack-user-vlan-id = 201
admin> set bridging-options packet-flows = ratelimit-wan
admin> set atm-options nailed-group = 251
admin> write -f

admin> new connection user-2
admin> set encapsulation-protocol = atm
admin> set ip-options ip-routing-enabled = no
admin> set bridging-options bridging-group = 7
admin> set bridging-options bridge = yes
admin> set bridging-options bridge-type = stacked-vlan
admin> set bridging-options vlan-stack-user-vlan-id = 202
admin> set bridging-options packet-flows = ratelimit-wan
admin> set atm-options nailed-group = 257
admin> write -f
```

## Configuring rate limiting and traffic shaping for 1:1 VLAN bridging

Figure 5-9 shows a 1:1 vlan-circuit configuration in which a single user-side PVC is bridged to a single VLAN. In this sample configuration, the traffic transmitted over the VLAN interface will be rate-limited based on the flow classification. The traffic that was passed by the policer will then be shaped at 500kbps. Traffic transmitted in the downstream direction over the DSL interface is rate-limited as specified in the relevant PACKET-FLOWS profile.

Figure 5-9. 1:1 bridging VLAN with rate limiting and traffic shaping



To configure rate limiting and traffic shaping for a 1:1 VLAN bridging circuit, follow these steps:

- 1 Create a PACKET-FLOWS profile for rate-limiting traffic over the VLAN interface. For example:

```
admin> new packet-flows ratelimit-vlan
admin> set flow 1 packet-classifier destination-ip-address = 224.0.0.0/4
admin> set flow 1 packet-classifier destination-ip-netmask = 240.0.0.0
admin> set flow 1 scheduling queue-priority = 7
admin> set flow 2 packet-classifier source-ip-address = 10.0.0.0/8
admin> set flow 2 packet-classifier source-ip-netmask = 255.0.0.0
admin> set flow 2 scheduling queue-priority = 6
admin> set flow 2 scheduling yellow-queue-priority = 2
admin> set flow 2 scheduling committed-information-rate = 100
admin> set flow 2 scheduling committed-burst-size = 2760
admin> set flow 2 scheduling excess-burst-size = 27600
admin> set flow 3 scheduling queue-priority = 1
admin> set flow 3 scheduling yellow-queue-priority = 0
admin> set flow 3 scheduling committed-information-rate = 20
admin> set flow 3 scheduling committed-burst-size = 3000
admin> set flow 3 scheduling excess-burst-size = 4500
admin> write -f
```
- 2 Configure a VLAN-ETHERNET profile for VLAN 1 with a scheduler of priority 1 and traffic shaping that transmits PDU traffic at 500 kbps, and attach the PACKET-FLOWS profile. For example:

```
admin> new vlan-ethernet
admin> set interface-address physical shelf = shelf-1
admin> set interface-address physical slot = first-control-module
```

```
admin> set interface-address physical item-number = 2
admin> set interface-address logical-item = 1
admin> set enabled = yes
admin> set bridging-options bridging-group = 7
admin> set bridging-options bridge = yes
admin> set bridging-options bridge-type = vlan-circuit
admin> set bridging-options packet-flows = ratelimit-vlan
admin> set scheduling priority = 1
admin> set scheduling bandwidth = 500
admin> write -f
```

- 3 Create a PACKET-FLOWS profile for rate-limiting traffic over the DSL interface of the VLAN circuit. For example:

```
admin> new packet-flows ratelimit-wan
admin> set flow 1 packet-classifier destination-ip-address = 224.0.0.0/4
admin> set flow 1 packet-classifier destination-ip-netmask = 240.0.0.0
admin> set flow 1 scheduling queue-priority = 7
admin> set flow 2 packet-classifier source-ip-address = 10.0.0.0/8
admin> set flow 2 packet-classifier source-ip-netmask = 255.0.0.0
admin> set flow 2 scheduling queue-priority = 6
admin> set flow 2 scheduling yellow-queue-priority = 2
admin> set flow 2 scheduling committed-information-rate = 100
admin> set flow 2 scheduling committed-burst-size = 2760
admin> set flow 2 scheduling excess-burst-size = 27600
admin> set flow 3 scheduling queue-priority = 1
admin> set flow 3 scheduling yellow-queue-priority = 0
admin> set flow 3 scheduling committed-information-rate = 20
admin> set flow 3 scheduling committed-burst-size = 3000
admin> set flow 3 scheduling excess-burst-size = 4500
admin> write -f
```

- 4 Create the CONNECTION profile for the ATM PVC and attach the PACKET-FLOWS profiles for rate limiting. For example:

```
admin> new connection user-1
admin> set encapsulation-protocol = atm
admin> set ip-options ip-routing-enabled = no
admin> set bridging-options bridging-group = 7
admin> set bridging-options bridge = yes
admin> set bridging-options bridge-type = vlan-circuit
admin> set bridging-options packet-flows = ratelimit-wan
admin> set atm-options nailed-group = 251
admin> write -f
```

## **Configuring a PPPoE connection with and without line-rate tracking**

The following commands configure a PPPoE connection:

```
admin> new connection
admin> set station = pppoe-1
admin> set active = yes
admin> set encapsulation-protocol = ppp
admin> set ip-options remote-address = 3.3.3.4/8
admin> set ip-options local-address = 3.3.3.3/8
admin> set session-options idle-timer = 0
admin> set session-options ts-idle-timer = 0
admin> set telco-options call-type = off
admin> set ppp-options send-auth-mode = pap-ppp-auth
admin> set ppp-options recv-password = ascend
admin> set ppp-options link-compression = none
admin> set ppp-options lqm = yes
admin> set mp-options enabled = no
admin> set pppoe-options pppoe = yes
admin> set pppoa-options pppoa = no
admin> write -f
```

The following commands create an ATM-QOS profile with a PCR of 12Mbps without the DSL line-rate tracking feature:

```
admin> new atm-qos
admin> set contract-name = dn-q
admin> set peak-rate-kbits-per-sec = 12000
admin> set peak-cell-rate-cells-per-sec = 28301
admin> set sustainable-rate-kbits-per-sec = 12000
admin> set sustainable-cell-rate-cells-per-sec = 28301
admin> write -f
```

The following commands apply the ATM-QOS profile to downstream traffic on the terminating ATM circuit profile:

```
admin> new connection
admin> set station = atm-ckt
admin> set active = yes
admin> set atm-options vci = 36
admin> set atm-options nailed-group = 7
admin> set atm-connect-options vci = 36
admin> set atm-connect-options nailed-group = 2097
admin> set atm-qos-options usr-dn-stream-contract = dn-q
admin> write -f
```

After the PPPoE session has been established and downstream traffic is being sent, if the underlying DSL line-rate adapts below the configured PCR of 12Mbps to 10Mbps (for example), the CM NP continues to shape traffic at the expected QoS of 12Mbps. As a result, 2Mbps will be discarded on the LIM. The “TX queue drop packets” value in the following command output shows that the LIM is discarding cells:

```
admin> open 1 1
ep-hs-gs-48-1/1> stats error ffwBridge 7 0 36
DSL AAL5 channel statistics
 Rx Aborted frames : 0
 Rx CRC error : 0
 Rx underflow error : 0
 Rx max SDU error : 0
 Rx IW underrun error : 0
 Rx IW max MRU error : 0
 Rx CPI error : 0
 Rx RAS timeout error : 0
DSL flow statistics
 FBP drop packets : 0
 MTU drop packets : 0
 TTL drop packets : 0
 TX queue drop packets : 9876
 Denied packets : 0
 Group filtered packets : 0
 GTP bad header : 0

Sonet AAL5 channel statistics
 Rx Aborted frames : 0
 Rx CRC error : 0
 Rx underflow error : 0
 Rx max SDU error : 0
 Rx IW underrun error : 0
 Rx IW max MRU error : 0
 Rx CPI error : 0
 Rx RAS timeout error : 0
Sonet flow statistics
 FBP drop packets : 0
 MTU drop packets : 0
 TTL drop packets : 0
 TX queue drop packets : 0
 Denied packets : 0
 Group filtered packets : 0
 GTP bad header : 0

ep-hs-gs-48-1/1> quit
```

The following commands modify the ATM-QOS profile to enable the DSL line-rate tracking feature:

```
admin> read atm-qos dn-q
admin> set track-underlying-dsl = yes
admin> write -f
```

After the system has bounced the PPPoE session, the CM NP begins to discard packets on the basis of CoS priorities, and no cell discards will be seen on the LIM.

For example, the following commands show CM PDU discards incrementing:

```
admin> stats traffic np lport 0
lportID blkRemain tttPass tttDiscard pduPass pduDiscard Thresh
 0 199 5643550 14712 2821808 7356 8170

admin> stats traffic np lport 0
lportID blkRemain tttPass tttDiscard pduPass pduDiscard Thresh
 0 200 5652972 15372 2826518 7686 8170
```

### **Configuring a BIR connection with line-rate tracking**

The following commands create an ATM-QOS profile with a PCR of 12Mbps and enabled DSL line-rate tracking:

```
admin> new atm-qos
admin> set contract-name = dn-q
admin> set peak-rate-kbits-per-sec = 12000
admin> set peak-cell-rate-cells-per-sec = 28301
admin> set sustainable-rate-kbits-per-sec = 12000
admin> set sustainable-cell-rate-cells-per-sec = 28301
admin> set track-underlying-dsl = yes
admin> write -f
```

The following commands configure a BIR connection and apply the ATM-QOS profile:

```
admin> new connection
admin> set station = bir
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 12.12.12.3/24
admin> set ip-options netmask-remote = 255.255.255.0
admin> set ip-options local-address = 12.12.12.2/32
admin> set ip-options netmask-local = 255.255.255.255
admin> set atm-options vpi = 10
admin> set atm-options nailed-group = 7
admin> set atm-qos-options usr-dn-stream-contract = dn-q
admin> set bir-options enable = yes
admin> write -f
```

## **Administrative tools for monitoring IP QoS**

The `stats` command, described in detail in the *Stinger Reference*, provides options for gathering and displaying traffic statistics related to QoS. This section describes only the commands related to monitoring IP QoS.

Following is the syntax of the `stats` command in the context of IP QoS monitoring:

```
stats cmd np ipqos ifnum [up | down]
```

The `cmd` argument specifies one of the following actions to be performed on the specified interface (`ifnum`):

- `enable` (Begin monitoring QoS on the specified interface.)
- `disable` (Stop monitoring QoS on the specified interface.)
- `traffic` (Display traffic statistics on the specified interface.)
- `rate` (Display traffic statistics each second. Entering the same command again causes the frequent updates to stop.)
- `error` (Display error statistics on the specified interface.)
- `clear` (Clear statistics on the specified interface.)

The `ifnum` argument must be obtained via the `ifmgr -d` command, which is available only in the debug environment. For details about `ifmgr -d`, see “ifmgr” on page A-14.

The `[ up | down ]` argument applies only to `vlan-circuit` bridged connections. For these connections the `ifnum` argument is always that of the DSL WAN interface number, and the `up` or `down` argument is used to enable QoS monitoring upstream (on the VLAN interface), or downstream (on the WAN interface).

For example, the following command starts QoS statistics monitoring on interface 15:

```
admin> stats enable np ipqos 15
```

and the following command stops monitoring on that interface:

```
admin> stats disable np ipqos 15
```

### **Example of monitoring routed traffic onto Gigabit Ethernet**

This example shows IP QoS monitoring on a Gigabit Ethernet interface as routed traffic is forwarded onto Ethernet.

#### **Creating a terminating routed connection**

The following commands configure a routed connection:

```
admin> new connection routed-test
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 4.4.4.2/8
admin> set ip-options local-address = 4.4.4.1/8
admin> set atm-options vci = 70
admin> set atm-options nailed-group = 101
admin> write -f
```

## Applying a PACKET-FLOWS profile to the Ethernet IP interface

The following commands configure a PACKET-FLOWS profile and apply it to the IP-INTERFACE profile:

```
admin> new packet-flows
admin> set name = ipflows
admin> set flow 1 packet-classifier source-ip-address = 9.0.0.0/8
admin> set flow 1 packet-classifier source-ip-netmask = 255.0.0.0
admin> set flow 1 scheduling queue-priority = 6
admin> set flow 2 packet-classifier source-ip-address = 8.0.0.0/8
admin> set flow 2 packet-classifier source-ip-netmask = 255.0.0.0
admin> set flow 2 scheduling queue-priority = 5
admin> set flow 3 packet-classifier source-ip-address = 7.0.0.0/8
admin> set flow 3 packet-classifier source-ip-netmask = 255.0.0.0
admin> set flow 3 scheduling queue-priority = 4
admin> write -f
admin> read ip-interface { { 1 8 2 } 0 }
admin> set ip-address = 7.7.7.1/8
admin> set packet-flows = ipflows
admin> write -f
```

## Obtaining the Ethernet interface number

The `ifmgr -d` command, available in the debug environment, obtains the Gigabit Ethernet interface number (interface 1 in the output below). For details, see Appendix A, “IP Control Module Diagnostics.”

```
admin> ifmgr -d
bif slot sif u m p ifname host-name remote-addr local-addr

000 1:08 000 * ie0 - 0.0.0.0/32 110.110.110.72/32
001 1:08 001 * ie1 - 0.0.0.0/32 7.7.7.1/32
006 1:08 006 * lo0 - 0.0.0.0/32 127.0.0.1/32
007 0:00 000 * rj0 - 0.0.0.0/32 127.0.0.2/32
008 0:00 000 * bh0 - 0.0.0.0/32 127.0.0.3/32
009 1:08 000 * wanabe - 0.0.0.0/32 127.0.0.3/32
010 0:00 000 * local - 0.0.0.0/32 127.0.0.1/32
011 0:00 000 * mcast - 0.0.0.0/32 224.0.0.0/32
012 0:00 000 - tunnel0 - 0.0.0.0/32 127.0.0.5/32
013 0:00 000 * vr0_main - 0.0.0.0/32 127.0.0.4/32
014 0:00 000 - sip0 - 0.0.0.0/32 0.0.0.0/32
017 1:03 002 * p wan17 x 4.4.4.2/32 4.4.4.1/32
<end>
```

## Enabling monitoring on the Ethernet IP interface

The following command shows that the CoS-specific QoS counters are incremented on interface 1 (the Gigabit Ethernet port):



```

admin> stats enable np ipqos 1
IP QoS monitoring start for IF 1
IP QoS profile: pbitmark
CLASSIFICATION & FLOW
 Default classification :provided
 Number of classified DID :0
 Number of class rules :5
 Number of multicast IP flow :0
 Number of unicast IP flow :0
PRIORITIES
 Number of priority :4
 Lowest priority :0
 Mapping priority/COS :P[0]=5 P[1]=x P[2]=x P[3]=x P[4]=4 P[5]=3
P[6]=2 P[7]=x
SCHEDULING
 Line Rate :0
 Data queue Id :256
 Scheduler ID :3
 PDU type :packet
 Number of COS queues :5
 Scheduling priority :0
 Peak Rate (Kbps) :0
 TS parameters :00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 TM parameters :06 0f 7d 00 05 dc 00 00 00 00 00 00 00 00 00 00

admin> stats traffic np ipqos 1
FPP counters for IF 1

Priority	COS	Counter
drp	-	0
exp	1	0
6	2	179
5	3	0
4	4	0
0	5	0

TM counters for IF 1

Priority	COS	MTU drop	Queue drop	Sched drop	Port drop
exp	1	0	0	0	0
6	2	0	0	0	0
5	3	0	0	0	0
4	4	0	0	0	0
0	5	0	0	0	0


```

### **Example of monitoring bridged VLAN traffic (transparent bridging)**

This example shows IP QoS monitoring on both a DSL WAN interface and a VLAN interface, for an N:1 bridged VLAN configuration.

#### **Creating a bridged VLAN interface and PACKET-FLOWS profile**

The following commands create a bridged VLAN interface and a PACKET-FLOWS profile to monitor its traffic:

```

admin> new bridge-group
admin> set enable = yes

```

```
admin> set bridging-group = 572
admin> set port-block-enabled = no
admin> write -f
admin> new vlan-ethernet { { 1 8 2 } 572 }
admin> set enabled = yes
admin> set bridging-options bridging-group = 572
admin> set bridging-options bridge = yes
admin> set bridging-options bridge-type = transparent-bridging
admin> set bridging-options packet-flows = pbitmark
admin> write -f
admin> new packet-flows pbitmark
admin> set marking-type = ethernet-priority
admin> set flow 1 packet-classifier source-ip-address = 9.0.0.0/8
admin> set flow 1 packet-classifier source-ip-netmask = 255.0.0.0
admin> set flow 1 scheduling queue-priority = 6
admin> set flow 2 packet-classifier source-ip-address = 8.0.0.0/8
admin> set flow 2 packet-classifier source-ip-netmask = 255.0.0.0
admin> set flow 2 scheduling queue-priority = 5
admin> set flow 3 packet-classifier source-ip-address = 7.0.0.0/8
admin> set flow 3 packet-classifier source-ip-netmask = 255.0.0.0
admin> set flow 3 scheduling queue-priority = 4
admin> set global-packet-marking 5 ethernet-priority = 5
admin> set global-packet-marking 6 ethernet-priority = 6
admin> set global-packet-marking 7 ethernet-priority = 7
admin> write -f
```

### Creating a bridged subscriber interface and PACKET-FLOWS profile

The following commands configure a bridged VLAN subscriber interface and a PACKET-FLOWS profile for its traffic flow:

```
admin> new connection bridged-vlan-test
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options ip-routing-enabled = no
admin> set bridging-options bridging-group = 572
admin> set bridging-options bridge = yes
admin> set bridging-options bridge-type = transparent-bridging
admin> set bridging-options packet-flows = bridgesched
admin> set telco-options nailed-groups = 101
admin> set atm-options vci = 70
admin> set atm-options nailed-group = 101
```

```
admin> write -f
admin> new packet-flows bridgesched
admin> set flow 1 packet-classifier source-ip-address = 9.0.0.0/8
admin> set flow 1 packet-classifier source-ip-netmask = 255.0.0.0
admin> set flow 1 scheduling queue-priority = 6
admin> set flow 2 packet-classifier source-ip-address = 8.0.0.0/8
admin> set flow 2 packet-classifier source-ip-netmask = 255.0.0.0
admin> set flow 2 scheduling queue-priority = 5
admin> set flow 3 packet-classifier source-ip-address = 7.0.0.0/8
admin> set flow 3 packet-classifier source-ip-netmask = 255.0.0.0
admin> set flow 3 scheduling queue-priority = 4
admin> write -f
```

## Obtaining the interface numbers

The `ifmgr -d` command, available in the debug environment, obtains the VLAN and subscriber interface numbers (interfaces 16 and 15, respectively, in the output below). For details, see Appendix A, “IP Control Module Diagnostics.”

```
admin> ifmgr -d
bif slot sif u m p ifname host-name remote-addr local-addr

000 1:08 000 * ie0 - 0.0.0.0/32 110.110.110.72/32
001 1:08 001 * ie1 - 0.0.0.0/32 1.1.1.1/32
006 1:08 006 * lo0 - 0.0.0.0/32 127.0.0.1/32
007 0:00 000 * rj0 - 0.0.0.0/32 127.0.0.2/32
008 0:00 000 * bh0 - 0.0.0.0/32 127.0.0.3/32
009 1:08 000 * wanabe - 0.0.0.0/32 127.0.0.3/32
010 0:00 000 * local - 0.0.0.0/32 127.0.0.1/32
011 0:00 000 * mcast - 0.0.0.0/32 224.0.0.0/32
012 0:00 000 - tunnel0 - 0.0.0.0/32 127.0.0.5/32
013 0:00 000 * vr0_main - 0.0.0.0/32 127.0.0.4/32
014 0:00 000 - sip0 - 0.0.0.0/32 0.0.0.0/32
015 1:03 001 * p wan15 x 0.0.0.0/32 110.110.110.72/32
016 0:00 000 * vlan16 - 0.0.0.0/32 0.0.0.0/32
<end>
```

## Enabling monitoring for the bridged VLAN interface

The following commands enable IP QoS monitoring on the bridged VLAN interface (interface 16):

```
admin> stats enable np ipqos 16
IP QOS monitoring start for IF 16
IP QOS profile: pbitmark
CLASSIFICATION & FLOW
 Default classification :provided
 Number of classified DID :3
 Number of class rules :5
 Number of multicast IP flow :0
 Number of unicast IP flow :0
```

```

PRIORITIES
 Number of priority :4
 Lowest priority :0
 Mapping priority/COS :P[0]=5 P[1]=x P[2]=x P[3]=x P[4]=4 P[5]=3
P[6]=2 P[7]=x
SCHEDULING
 Line Rate :0
 Data queue Id :274
 Scheduler ID :3
 PDU type :packet
 Number of COS queues :5
 Scheduling priority :0
 Peak Rate (Kbps) :0
 TS parameters :00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 TM parameters :06 0f 7d 00 05 dc 00 01 00 00 00 00 00 00 00 00

admin> stats traffic np ipqos 16
FPP counters for IF 16
Priority COS Counter
 drp - 0
 exp 1 0
 6 2 143
 5 3 0
 4 4 0
 0 5 0
TM counters for IF 16
Priority COS MTU drop Queue drop Sched drop Port drop
exp 1 0 0 0 0
6 2 0 0 0 0
5 3 0 0 0 0
4 4 0 0 0 0
0 5 0 0 0 0

```

## Enabling monitoring for the bridged subscriber interface

The following commands disable monitoring on the VLAN interface (16) and then enable monitoring on the bridged VLAN subscriber interface (interface 15) with a PACKET-FLOWS profile attached.

```

admin> stats disable np ipqos 16
admin> stats enable np ipqos 15
IP QOS monitoring start for IF 15
IP QOS profile: bridgesched
CLASSIFICATION & FLOW
 Default classification :provided
 Number of classified DID :3
 Number of class rules :5
 Number of multicast IP flow :0
 Number of unicast IP flow :0
PRIORITIES
 Number of priority :4
 Lowest priority :0
 Mapping priority/COS :P[0]=5 P[1]=x P[2]=x P[3]=x P[4]=4 P[5]=3
P[6]=2 P[7]=x

```

```
SCHEDULING
 Line Rate :0
 Data queue Id :262
 Scheduler ID :3
 PDU type :cell
 Number of COS queues :5
 ATM service category :UBR
 PCR (cell per sec) :0
 TS parameters :00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 TM parameters :06 4b 00 c8 00 c8 00 02 00 00 00 00 00 00 00 00

admin> stats traffic np ipqos 15
FPP counters for IF 15
Priority COS Counter
 drp - 0
 exp 1 0
 6 2 173
 5 3 0
 4 4 0
 0 5 0

TM counters for IF 15
Priority COS MTU drop Queue drop Sched drop Port drop
exp 1 0 0 0 0
6 2 0 0 0 0
5 3 0 0 0 0
4 4 0 0 0 0
0 5 0 0 0 0
```

### **Example of monitoring 1:1 VLAN bridged traffic**

For vlan-circuit bridged connections, you monitor QoS on the WAN interface only, and specify the direction of traffic (up or down).

This example shows IP QoS monitoring on the DSL WAN interface in the upstream direction for a 1:1 (vlan-circuit) bridged VLAN.

### **Creating the subscriber-side profiles**

The following commands configure the subscriber-side CONNECTION profile and a PACKET-FLOWS profile to monitor its traffic:

```
admin> new connection vlancir1
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options ip-routing-enabled = no
admin> set bridging-options bridging-group = 572
admin> set bridging-options bridge = yes
admin> set bridging-options packet-flows = bridgesched
admin> set atm-options vci = 70
admin> set atm-options nailed-group = 101
admin> write -f
```

```
admin> new packet-flows bridgesched
admin> set flow 1 packet-classifier source-ip-address = 9.0.0.0/8
admin> set flow 1 packet-classifier source-ip-netmask = 255.0.0.0
admin> set flow 1 scheduling queue-priority = 6
admin> set flow 2 packet-classifier source-ip-address = 8.0.0.0/8
admin> set flow 2 packet-classifier source-ip-netmask = 255.0.0.0
admin> set flow 2 scheduling queue-priority = 5
admin> set flow 3 packet-classifier source-ip-address = 7.0.0.0/8
admin> set flow 3 packet-classifier source-ip-netmask = 255.0.0.0
admin> set flow 3 scheduling queue-priority = 4
admin> write -f
```

## Creating the VLAN-side profiles

The next set of commands configure the VLAN-ETHERNET profile for the 1:1 bridged VLAN, and a PACKET-FLOWS profile to monitor its traffic:

```
admin> new vlan-ethernet { { 1 8 2 } 572 }
admin> set enabled = yes
admin> set bridging-options bridge = yes
admin> set bridging-options packet-flows = pbitmark
admin> write -f
admin> new packet-flows pbitmark
admin> set marking-type = ethernet-priority
admin> set flow 1 packet-classifier source-ip-address = 9.0.0.0/8
admin> set flow 1 packet-classifier source-ip-netmask = 255.0.0.0
admin> set flow 1 scheduling queue-priority = 6
admin> set flow 2 packet-classifier source-ip-address = 8.0.0.0/8
admin> set flow 2 packet-classifier source-ip-netmask = 255.0.0.0
admin> set flow 2 scheduling queue-priority = 5
admin> set flow 3 packet-classifier source-ip-address = 7.0.0.0/8
admin> set flow 3 packet-classifier source-ip-netmask = 255.0.0.0
admin> set flow 3 scheduling queue-priority = 4
admin> set global-packet-marking 5 ethernet-priority = 5
admin> set global-packet-marking 6 ethernet-priority = 6
admin> set global-packet-marking 7 ethernet-priority = 7
admin> write -f
```

## Obtaining the interface numbers

The `ifmgr -d` command, available in the debug environment, obtains the subscriber interface number (interface 15 in the output below). For details, see Appendix A, “IP Control Module Diagnostics.”

```
admin> ifmgr -d
```

| bif | slot | sif | u | m | p | ifname   | host-name | remote-addr | local-addr        |
|-----|------|-----|---|---|---|----------|-----------|-------------|-------------------|
| 000 | 1:08 | 000 | * |   |   | ie0      | -         | 0.0.0.0/32  | 110.110.110.72/32 |
| 001 | 1:08 | 001 | * |   |   | ie1      | -         | 0.0.0.0/32  | 1.1.1.1/32        |
| 006 | 1:08 | 006 | * |   |   | lo0      | -         | 0.0.0.0/32  | 127.0.0.1/32      |
| 007 | 0:00 | 000 | * |   |   | rj0      | -         | 0.0.0.0/32  | 127.0.0.2/32      |
| 008 | 0:00 | 000 | * |   |   | bh0      | -         | 0.0.0.0/32  | 127.0.0.3/32      |
| 009 | 1:08 | 000 | * |   |   | wanabe   | -         | 0.0.0.0/32  | 127.0.0.3/32      |
| 010 | 0:00 | 000 | * |   |   | local    | -         | 0.0.0.0/32  | 127.0.0.1/32      |
| 011 | 0:00 | 000 | * |   |   | mcast    | -         | 0.0.0.0/32  | 224.0.0.0/32      |
| 012 | 0:00 | 000 | - |   |   | tunnel0  | -         | 0.0.0.0/32  | 127.0.0.5/32      |
| 013 | 0:00 | 000 | * |   |   | vr0_main | -         | 0.0.0.0/32  | 127.0.0.4/32      |
| 014 | 0:00 | 000 | - |   |   | sip0     | -         | 0.0.0.0/32  | 0.0.0.0/32        |
| 015 | 1:03 | 001 | * |   | p | wan15    | vlancir1  | 0.0.0.0/32  | 110.110.110.72/32 |
| 016 | 0:00 | 000 | * |   |   | vlan16   | -         | 0.0.0.0/32  | 0.0.0.0/32        |

<end>

## Enabling monitoring in the upstream direction

The following commands enable IP QoS monitoring in upstream direction and show that the QoS counters are incremented.

```
admin> stats enable np ipqos 15 up
IP QOS monitoring start for IF 15 flow direction - upstream
IP QOS profile: pbitmark
QoS monitoring direction: upstream
CLASSIFICATION & FLOW
 Default classification :provided
 Number of classified DID :3
 Number of class rules :5
 Number of multicast IP flow :0
 Number of unicast IP flow :0
PRIORITIES
 Number of priority :4
 Lowest priority :0
 Mapping priority/COS :P[0]=5 P[1]=x P[2]=x P[3]=x P[4]=4 P[5]=3
P[6]=2 P[7]=x
SCHEDULING
 Line Rate :0
 Data queue Id :274
 Scheduler ID :3
 PDU type :packet
 Number of COS queues :5
 Scheduling priority :0
 Peak Rate (Kbps) :0
 TS parameters :00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 TM parameters :06 0f 7d 00 05 dc 00 01 00 00 00 00 00 00 00 00
admin> stats traffic np ipqos 15
QoS monitoring direction: upstream
FPP counters for IF 15
Priority COS Counter
drp - 0
exp 1 0
```

|   |   |     |
|---|---|-----|
| 6 | 2 | 150 |
| 5 | 3 | 0   |
| 4 | 4 | 0   |
| 0 | 5 | 0   |

TM counters for IF 15

| Priority | COS | MTU drop | Queue drop | Sched drop | Port drop |
|----------|-----|----------|------------|------------|-----------|
| exp      | 1   | 0        | 0          | 0          | 0         |
| 6        | 2   | 0        | 0          | 0          | 0         |
| 5        | 3   | 0        | 0          | 0          | 0         |
| 4        | 4   | 0        | 0          | 0          | 0         |
| 0        | 5   | 0        | 0          | 0          | 0         |

### Enabling monitoring in the downstream direction

The following commands enable IP QoS monitoring in downstream direction and show that the QoS counters are incremented.

```
admin> stats enable np ipqos 15 down
IP QOS monitoring start for IF 15 flow direction - downstream
IP QOS profile: bridgesched
QoS monitoring direction: downstream
CLASSIFICATION & FLOW
 Default classification :provided
 Number of classified DID :3
 Number of class rules :5
 Number of multicast IP flow :0
 Number of unicast IP flow :0
PRIORITIES
 Number of priority :4
 Lowest priority :0
 Mapping priority/COS :P[0]=5 P[1]=x P[2]=x P[3]=x P[4]=4 P[5]=3
P[6]=2 P[7]=x
SCHEDULING
 Line Rate :0
 Data queue Id :262
 Scheduler ID :3
 PDU type :cell
 Number of COS queues :5
 ATM service category :UBR
 PCR (cell per sec) :0
 TS parameters :00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 TM parameters :06 4b 00 c8 00 c8 00 02 00 00 00 00 00 00 00 00
admin> stats traffic np ipqos 15
QoS monitoring direction: downstream
FPP counters for IF 15
Priority COS Counter
drp - 0
exp 1 0
6 2 139
5 3 0
4 4 0
0 5 0
TM counters for IF 15
```



| Priority | COS | MTU drop | Queue drop | Sched drop | Port drop |
|----------|-----|----------|------------|------------|-----------|
| exp      | 1   | 0        | 0          | 0          | 0         |
| 6        | 2   | 0        | 0          | 0          | 0         |
| 5        | 3   | 0        | 0          | 0          | 0         |
| 4        | 4   | 0        | 0          | 0          | 0         |
| 0        | 5   | 0        | 0          | 0          | 0         |

### **Example of monitoring stacked VLAN bridged traffic**

For stacked-vlan bridged connections, you can enable QoS monitoring on an NSP VLAN interface or on a stacked (pre-tagged or untagged) WAN interface.

#### **Enabling QoS monitoring on an NSP VLAN interface**

For stacked VLAN connections, you enable or disable QoS monitoring on an NSP VLAN interface by using the following commands:

```
admin> info np conn
Connection Handles
if slot act cPort dPort type uctl udata qos dctl ddata qos
1 8 Y 000000 000000 GE 0/870 0/870 default 0/870 0/870 default
21 8 Y 000000 000000 GE 0/870 0/870 default 0/870 0/870 default
15 8 Y 000000 000008 CM 0/871 0/871 default 0/871 0/871 default
Total 3 Connections <end>
```

```
admin> stats enable np ipqos 21
IP QOS monitoring start for IF 21
IP QOS profile: ratelimit-vlan
CLASSIFICATION & FLOW
 Default classification :provided
 Number of classified DID :0
 Number of class rules :0
 Number of multicast IP flow :0
 Number of unicast IP flow :2
 Unicast IP flow 2
 Configuration ID :2
 Flow ID :10
 CIR (Kbps) :100
 CBS (bytes) :2760
 EBS (bytes) :27600
 Green Priority/Cos ID :6/3
 Yellow Priority/Cos ID :2/4
 Red Priority/Cos ID :-1/0
 Unicast IP flow 3
 Configuration ID :3
 Flow ID :11
 CIR (Kbps) :20
 CBS (bytes) :3000
 EBS (bytes) :4500
 Green Priority/Cos ID :1/5
 Yellow Priority/Cos ID :0/6
 Red Priority/Cos ID :-1/0
```

```

PRIORITIES
 Number of priority :5
 Lowest priority :0
 Mapping priority/COS :P[0]=6 P[1]=5 P[2]=4 P[3]=x P[4]=x P[5]=x
P[6]=3 P[7]=2
SCHEDULING
 Line Rate :0
 Data queue Id :263
 Scheduler ID :3
 PDU type :packet
 Number of COS queues :6
 Scheduling priority :0
 Peak Rate (Kbps) :0
 TS parameters :00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 TM parameters :06 0f 7d 00 05 dc 00 00 00 00 00 00 00 00 00 00

```

The following command disabled QoS monitoring on the NSP VLAN interface:

```
admin> stats disable np ipqos 21
```

## Enabling QoS monitoring on a stacked VLAN WAN connection

The following command enables QoS monitoring on the stacked-VLAN WAN interface:

```

admin> stats enable np ipqos 15
IP QOS monitoring start for IF 15
IP QOS profile: ratelimit-wan
CLASSIFICATION & FLOW
 Default classification :provided
 Number of classified DID :4
 Number of class rules :4
 Number of multicast IP flow :0
 Number of unicast IP flow :2
 Unicast IP flow 2
 Configuration ID :2
 Flow ID :10
 CIR (Kbps) :100
 CBS (bytes) :2760
 EBS (bytes) :27600
 Green Priority/Cos ID :6/3
 Yellow Priority/Cos ID :2/4
 Red Priority/Cos ID : -1/0
 Unicast IP flow 3
 Configuration ID :3
 Flow ID :11
 CIR (Kbps) :20
 CBS (bytes) :3000
 EBS (bytes) :4500
 Green Priority/Cos ID :1/5
 Yellow Priority/Cos ID :0/6
 Red Priority/Cos ID : -1/0

```

```

PRIORITIES
 Number of priority :5
 Lowest priority :0
 Mapping priority/COS :P[0]=6 P[1]=5 P[2]=4 P[3]=x P[4]=x P[5]=x
 P[6]=3 P[7]=2

```

### **Example of displaying a connection's traffic shaping parameters**

Both Stinger IP control modules support the debug-level `info np` command. For example, the `info np conn` command displays queue IDs on the CM, which is needed to view traffic shaping parameters on the IP2000 module.

On the IP2100 module, the `agsh` command displays both the queue ID and traffic shaping parameters.

For details about the debug environment, see Appendix A, "IP Control Module Diagnostics."

### **Displaying a connection's traffic shaping parameters on the IP2000**

To display traffic shaping parameters on the IP2000 control module, complete the following steps in the debug environment.

- 1 Use the `info np conn` command to display the queue IDs for a connection.

The queue IDs are listed under the Principal DID list near the bottom of the command output. For example, the following command displays queue IDs related to a PPPoE session:

```

super> info np conn 1 1 18
Connection Handles
if sh/sl act cPort dPort type uctl udata dctl ddata
18 1/1 Y 000000 000000 LIM 0/149 0/148 0/149 0/148
Encaps: PPPoE - MTU=1492 (50)

```

Control & Data Service Queues:

-----

|            |     |          |        |
|------------|-----|----------|--------|
| ControlQID | 272 | Lport    | 0x0000 |
| Scheduler  | CBR | CosQueue | 1      |
| DataQID    | 276 | Lport    | 0x0000 |
| Scheduler  | CBR | CosQueue | 3      |

PDU Ids:

-----

| PduID  | port   | vpi/vci | Tree |
|--------|--------|---------|------|
| 852099 | 0x0200 | 0/149   | 3072 |
| 327812 | 0x0200 | 0/148   | 3072 |

VLAN Configuration : Vlan is not enabled

Packet Flow configuration: No Packet Flow is attached

Multicast Configuration : Multicast is not enabled

Bridge Configuration : Bridging is not enabled

Filter Configuration : No Filters applied

TCP Estab Outgoing Filter: Not enabled

QOS Configuration :  
     Downstream : dn-q  
     Upstream : default

Principal DID List:

| Did   | Type    | QId | cosId | Cloned |
|-------|---------|-----|-------|--------|
| 2006  | DATA_EX | 277 | 1     | No     |
| 55156 | UNICAST | 279 | 3     | No     |

Layer 2 DID List: Empty

## Displaying a connection's traffic shaping parameters on the IP2100

To display traffic shaping parameters on the IP2100 control module, complete the following steps in the debug environment.

- 1 Use the `info np conn` command to display the queue IDs for a connection.

The queue IDs are listed under the Principal DID list near the bottom of the command output. For example, the following command displays queue IDs related to a BIR connection:

```
super> info np conn 1 1 17
Connection Handles
if sh/sl act cPort
type uctl udata dctl ddata
17 1/1 Y 000000 000000 LIM 0/149 0/148 0/149 0/148
Encaps: BIR - MTU=1560 (50)
```

Control & Data Service Queues:

```

ControlQID 276 Lport 0x0000
Scheduler CBR CosQueue 1

DataQID 272 Lport 0x0000
Scheduler CBR CosQueue 3
```

PDU Ids:

```

PduID port vpi/vci Tree
065667 0x0200 0/149 3072
196740 0x0200 0/148 3072
```

VLAN Configuration : Vlan is not enabled

Packet Flow configuration: No Packet Flow is attached

Multicast Configuration : Multicast is not enabled

Bridge Configuration : Bridging is not enabled

Filter Configuration : No Filters applied

TCP Estab Outgoing Filter: Not enabled

QOS Configuration :  
Downstream : dn-q  
Upstream : default

Principal DID List:

| Did   | Type    | QId | cosId | Cloned |
|-------|---------|-----|-------|--------|
| 2004  | CNTL_EX | 277 | 1     | No     |
| 2005  | DATA_EX | 273 | 1     | No     |
| 55158 | UNICAST | 275 | 3     | No     |
| 55159 | ARP     | 275 | 3     | No     |

Layer 2 DID List: Empty

- Using the queue IDs displayed at the bottom of the output immediately above, use the agsh npTTPS command to display traffic shaping parameters on the CM NP.

For example:

```
super> agsh npTTPS -i 0 275
parameter byte: 00:01:02:03:04:05:06:07:08:09:10:11:12:13:14:15:
Qid 275 -
TM addr 006037d0: 06:07:00:c8:00:c8:00:00:00:00:00:00:00:00:00:00
TS addr 004037d0: 00:00:00:91:21:45:03:be:00:92:01:00:00:00:00:00
```

## Limitations with the current software version

This section lists limitations that apply to the IP QoS implementation with this software version.

- A Stinger IP DSLAM can support classification among different multicast streams that are going to the same output interface. Thus, for example, two streams, one to the destination address 224.1.1.1 and the second to the destination address 224.1.1.2 but that go on the same DSL user connection can be classified to have 2 different priorities.

Prioritization between unicast and multicast flows on the same connection is also supported.

However, subclassification within a multicast flow cannot be done. Thus, for example, UDP packets on 224.1.1.1 and TCP packets going to 224.1.1.1 cannot be distinguished.

- Only 32 packet flows are supported per interface.

If 32 packet flows have been defined in the `packet-classifier` subprofile, and there is no default rule, the Stinger system creates an implicit default rule default.

- For fragmented IP packets, only the IP header is available for classification, which means that classification based on UDP/TCP source or destination ports cannot be used.
- Packets that are fragmented by the Stinger IP DSLAM always take the highest priority. Classifications rules for such packets are not applicable.
- IP QoS does not apply to L2TP connections.
- When a nonzero `source-port` or `destination-port` is specified in a `PACKET-FLOWS` profile, the `ip-protocol` value must be set to either 17 (UDP) or 6 (TCP).
- SNMP support is not provided with the current software version.
- The precision of the rate-limiting policer is about 1%.
- The minimum policing rate (the minimum CIR) is 17Kbps.
- The minimum scheduling rate for a logical interface on a Gigabit Ethernet port is 401Kbps.
- A given `PACKET-FLOWS` profile can have either `layer2-classifier` or IP-specific `packet-classifier` settings, but not both.

---

# ***L2TP Tunneling Configuration***

# **6**

|                                                             |      |
|-------------------------------------------------------------|------|
| Overview of L2TP tunneling .....                            | 6-2  |
| Overview of L2TP tunnel authentication.....                 | 6-3  |
| Configuring L2TP global options .....                       | 6-6  |
| Configuring client connections with PPP authentication..... | 6-11 |
| Configuring connection-based tunnel authentication .....    | 6-16 |
| Configuring server-based tunnel authentication .....        | 6-18 |
| Using tunnel assignment IDs .....                           | 6-22 |
| Administrative tools for monitoring L2TP .....              | 6-26 |
| Limitations with this software version.....                 | 6-27 |

Properly configured, a Stinger system with an IP control module can operate in Layer 2 Tunneling Protocol (L2TP) access concentrator (LAC) mode. In LAC mode, the system accepts incoming PPP client sessions (mobile clients) and routes PPP data traffic at wire speed to a remote L2TP network server (LNS).


Stinger systems support tunneling for both PPPoA and PPPoE mobile client connections. PPPoA and PPPoE connections use the same configuration parameters. The main difference between the two types of connection is that for PPPoE, the system supports several PPPoE sessions across the same port and VPI-VCI pair.

PPPoE sessions on the same port and VPI-VCI pair are handled individually. For example, one session can be terminated locally, and another session on the same port and VPI-VCI pair can be tunneled to an LNS via L2TP. Multiple PPPoE L2TP sessions on the same port and VPI-VCI pair can be tunneled to different LNS systems.

## Overview of L2TP tunneling

Figure 6-1 shows the elements of an L2TP tunnel through the IP-capable Stinger in LAC mode.

Figure 6-1. L2TP tunneling



A mobile client initiates a session, which causes the Stinger to initiate a tunnel to the LNS, establish the PPP session, and forward the data stream to the LNS across the IP network. The Stinger negotiates the PPP session, which can include user password authentication, and then negotiates the L2TP tunnel.

Tunnel authentication occurs before a tunnel is established between two end points. It is negotiated between the LAC and a tunnel server and is independent of user authentication. If tunnel authentication fails, all pending calls associated with the tunnel are dropped.

### Links to LNS servers

To take advantage of the IP control module wire-speed capability, the link between the LAC and the LNS should be across Gigabit Ethernet (with or without VLAN) or on a trunk port with AAL5 LLC encapsulation. The connection to the LNS is an IP link, which consists of a control link and zero or more data links. Both the control and data links are encapsulated in UDP.

The control link carries information that is used both to query whether the LNS can accept the current call and to establish a tunnel. L2TP implements a Hello mechanism by which the LAC and LNS verify that both are operational by sending each other a control message every minute or so. If the Hello message does not arrive within the configured hello-timer, the tunnel and all the tunneled connections are torn down.

Data links carry the client data, which consists of PPP frames. Each tunneled client connection supports one data link.

### Links to PPP clients

To enable establishment of PPPoA/PPPoE connections using IP routing, you must complete the following steps:

- 1 Configure the ANSWER-DEFAULTS profile to accept PPP session requests and require their authentication.
- 2 Configure a CONNECTION or RADIUS profile for the underlying ATM circuit required between the DSL interface on which the PPP connection will be established and the LIM internal interface.



- For PPPoE, the underlying circuit must use RFC 1483 LLC.
- 3 Configure a CONNECTION or RADIUS profile for each PPP connection.

All of these required steps are described in detail in “Configuring broadband RAS subscriber access” on page 4-85. Examples of LIM profiles and PPP client connections for L2TP tunneling are shown in this chapter.

## **L2TP-related network settings**

Table 6-1 describes network and RADIUS settings related to L2TP.

*Table 6-1. Network settings related to L2TP*

| Type of configuration     | Description                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Virtual routers           | When you associate a virtual router with an L2TP tunnel, all L2TP packets for that tunnel are sent only via that virtual router. For information about virtual routing, see Chapter 7, “Virtual Router Configuration.”                                                                                                                                                                    |
| UDP port for control link | When you change the setting of the <code>base-udp-port</code> parameter in the L2-TUNNEL-GLOBAL profile, a system reset is required for the L2TP subsystem to recognize the new UDP port number to use for the control link. The change does not take effect for existing active links until the link has been dropped.                                                                   |
| Stinger IP address        | Alcatel-Lucent recommends that you set the <code>system-ip-addr</code> parameter in the IP-GLOBAL profile of a Stinger unit that is operating as a LAC, particularly if the unit has multiple interfaces into the IP cloud. This setting is not required if RIP is enabled on the interfaces between the LAC and LNS, but it is recommended, because it helps to simplify configurations. |

## **L2TP-related RADIUS configuration**

RADIUS-authenticated sessions can make use of some L2TP tunnel features that are not supported in the local command-line interface. For example, the tunnel authentication RADIUS attribute-value pairs support tag fields, as described in RFC 2868, *RADIUS Attributes for Tunnel Protocol Support*.

Stinger units support L2TP tunnel accounting in accordance with RFC 2867, *RADIUS Accounting Modifications for Tunnel Protocol Support*.

For details, see the *TAOS RADIUS Guide and Reference*.

## **Overview of L2TP tunnel authentication**

Tunnel authentication occurs before a tunnel is established. It is negotiated between the LAC and LNS and is independent of user authentication. If tunnel authentication fails, all pending calls associated with the tunnel are dropped.

Tunnel authentication can use several types of identification, such as the LAC system name, a password, and a tunnel assignment ID. If RADIUS accounting is enabled, the system reports the IDs used for tunnel authentication in the Stop record.

## **System name used for tunnel authentication**

If tunnel authentication is enabled in the L2-TUNNEL-GLOBAL profile and the LAC is requesting a new tunnel, it looks for a system name to send to the LNS as follows:

- 1 Uses the `client-auth-id` if specified in the client's CONNECTION profile. If `client-auth-id` is not specified in the CONNECTION profile, the system goes on to the next alternative.
- 2 Uses the `client-auth-id` if specified in the TUNNEL-SERVER profile for the LNS. If `client-auth-id` is not specified in a TUNNEL-SERVER profile, the system goes on to the next alternative.
- 3 Uses the `l2tp-system-name` if specified in the L2-TUNNEL-GLOBAL profile. If `l2tp-system-name` is not specified in that profile, the system goes on to the next alternative.
- 4 Uses the `name` if specified in the SYSTEM profile. If the domain name is configured in the IP-GLOBAL profile, the specified system name is concatenated with the domain name. If `name` is not specified in the SYSTEM profile, the system goes on to the next alternative.
- 5 Sends the string `noname`.

## **How the system finds a matching tunnel**

Each LIM maintains one control channel (which uses one UDP port) for all tunnels used by DSL clients on the LIM. For that reason, the `base-udp-port` value in the L2-TUNNEL-GLOBAL profile is used together with the slot number to derive a base UDP port for each LIM. After the LAC has authenticated a PPP client's session request, it looks for an existing tunnel that matches both the tunnel server end point and the `client-auth-id` specified in the client's profile. If the LAC finds an established tunnel *on the same LIM* that matches these values, it uses the tunnel. If it does not find a matching tunnel on the same LIM, it initiates a tunnel request. This process can be used to create parallel L2TP tunnels by specifying different `client-auth-id` values in profiles.

If the client's profile specifies a hostname as the tunnel-server end point, the system must match both the hostname and the server's actual IP address to allow the client to use an established tunnel.

After comparing the tunnel transport address and, if specified in the client's profile, the tunnel server's hostname (`server-endpoint`) against existing tunnels, the system begins comparing the following optional parameters, in the order shown:

- `client-auth-id` specified in the client's profile and the `client-auth-id` used for existing tunnels
- `assignment-id` specified in the client's profile and the tunnel assignment ID of existing tunnels

If both the `client-auth-id` and the tunnel `assignment-id` values match an existing tunnel, the system uses that tunnel. A null value in any one of these fields in an existing tunnel matches only a null value in the corresponding parameter in the client profile. If the system does not find a matching tunnel entry, it initiates a new tunnel request.

If an `assignment-id` is not specified in the client's profile, the system attempts to match the client to an existing tunnel by using the `client-auth-id`.

If `client-auth-id` is not specified in the client's profile, the system attempts to match the client to an existing tunnel by using only the tunnel server's IP address (and hostname, if specified).

### ***Examples of how client-auth-id settings create parallel tunnels***

In this example, the LNS system's DNS hostname is `a.example.com` (a fully qualified domain name), which resolves to two IP addresses, `1.1.1.1` and `1.1.1.2`. The hostname `b.example.net` also resolves to the `1.1.1.1` address. Table 6-2 shows existing tunnels to the LNS, which were authenticated by different `client-auth-id` strings.

Table 6-2. Existing tunnels to the same LNS

| Address | client-auth-id | tunnel-server | Tunnel ID |
|---------|----------------|---------------|-----------|
| 1.1.1.1 | a1             | a.example.com | 102       |
| 1.1.1.1 | a2             | a.example.com | 103       |

Table 6-3 shows how the system matches the values in the clients' profiles as it receives session requests, and the resulting action the system takes in terms of using an existing tunnel or creating a new one.

Table 6-3. Tunnels created for clients based on profile settings

| Values used to match tunnel |                 |                | Resulting action          | Tunnel ID  |
|-----------------------------|-----------------|----------------|---------------------------|------------|
| Address:                    | client-auth-id: | tunnel-server: |                           |            |
| 1.1.1.1                     | a1              | a.example.com  | Reuse tunnel              | 102        |
| 1.1.1.1                     | a2              | a.example.com  | Reuse tunnel              | 103        |
| 1.1.1.1                     | b               | b.example.net  | Establish new tunnel      | 104        |
| 1.1.1.1                     | b               | a.example.com  | Establish new tunnel      | 105        |
| 1.1.1.1                     |                 | a.example.com  | Reuse tunnel <sup>1</sup> | 102 or 103 |
| 1.1.1.1                     | a2              | b.example.net  | Establish new tunnel      | 106        |
| 1.1.1.2                     | a1              | a.example.com  | Establish new tunnel      | 107        |

1. The client that does not supply a `client-auth-id` string matches the `tunnel-server` end point, so the existing tunnel to that end point (such as Tunnel ID 102) is reused.

### ***Examples of configuration errors causing multiple tunnels***

Configuration errors can lead to unintentional parallel tunnels to the same tunnel end point. For this reason, you must either use the `client-auth-id` setting for all user profiles to a particular LNS or decide *not* to use that setting for clients tunneling to that LNS.

For example, suppose your RADIUS users file contains the following two user profiles and tunnel server profile:

```
user1 Password = userpass
 Tunnel-Type = L2TP,
 Tunnel-Server-Endpoint = lns.example.com,
```

```
Tunnel-Client-Auth-ID = A-LAC,
...
user2 Password = userpass
Tunnel-Type = L2TP,
Tunnel-Server-Endpoint = lns.example.com,
...
lns.example.com User-Password = "", Service-Type = Dialout
Tunnel-Password = tunpass,
Tunnel-Client-Auth-ID = AllMyLACs
```

If user1 calls in first and establishes a tunnel, user2 can reuse that tunnel, as shown in Table 6-4.

Table 6-4. Tunnels created when user1 dials in first (error not detected)

| Values used to match tunnels                   |       |                 | Resulting action  | Tunnel ID |
|------------------------------------------------|-------|-----------------|-------------------|-----------|
| <b>Address:</b> client-auth-id: tunnel-server: |       |                 |                   |           |
| 2.2.2.2                                        | A-LAC | lns.example.com | Create new tunnel | 88        |
| 2.2.2.2                                        |       | lns.example.com | Reuse tunnel      | 88        |

However, if user2 calls in first and establishes a tunnel, the system obtains a system name for authentication from the TUNNEL-SERVER profile. When user1 dials in, the client is unable to reuse the tunnel, because the authentication names do not match. This result is shown in Table 6-5.

Table 6-5. Tunnels created when user2 dials in first (error shown)

| Values used to match tunnels                   |           |                 | Resulting action  | Tunnel ID |
|------------------------------------------------|-----------|-----------------|-------------------|-----------|
| <b>Address:</b> client-auth-id: tunnel-server: |           |                 |                   |           |
| 2.2.2.2                                        | AllMyLACs | lns.example.com | Create new tunnel | 40        |
| 2.2.2.2                                        | A-LAC     | lns.example.com | Create new tunnel | 42        |

## Configuring L2TP global options

To enable the Stinger to operate as LAC, you must enable L2TP and configure some tunneling parameters. The parameters in the L2-TUNNEL-GLOBAL profile are used globally for all LAC operations. They are not specific to one LNS.

With the exception of the `base-udp-port` parameter, which requires a system reset for a change to take effect, changes to the L2-TUNNEL-GLOBAL profile take effect as soon as possible after the profile is written.

### Enabling L2TP operations and authentication

Following are the parameters, shown with default values, that enable the system to operate as LAC:

```
[in L2-TUNNEL-GLOBAL]
l2tp-mode = disabled
l2tp-auth-enabled = no
l2tp-rx-window = 0
```

```
l2tp-system-name = ""
```

| Parameter         | Setting                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| l2tp-mode         | Enable/disable L2TP operations. With the default disabled value, L2TP is disabled. Set the parameter to lac to enable L2TP LAC operations.                                                                                                                                                                                                                                                                                                                                                        |
| l2tp-auth-enabled | Enable/disable L2TP tunnel authentication. By default, L2TP tunnel authentication is disabled. When you enable tunnel authentication, you can configure the level of authentication you need to verify each LNS before bringing up a control channel.                                                                                                                                                                                                                                             |
| l2tp-rx-window    | Advertised L2TP receive window size for data channels, from 0 to 63. The default, 0 (zero), specifies that the LAC asks for no flow control for inbound L2TP payloads.<br><br><b>Note</b> Alcatel-Lucent strongly recommends a setting of zero for this parameter. A nonzero setting indicates that the LAC should perform sequencing of the data channel, which is not supported on Stinger IP DSLAMs. For additional information, see “Impact of L2TP data sequencing limitation” on page 6-28. |
| l2tp-system-name  | Name (up to 31 characters) that can be sent to the LNS during tunnel authentication.                                                                                                                                                                                                                                                                                                                                                                                                              |

### **Setting L2TP timers and other variables**

Following are the parameters, shown with default values, for configuring timers and other variables that affect tunnel retry attempts and other control message variables:

```
[in L2-TUNNEL-GLOBAL:l2tp-config]
first-retry-timer = 1000
retry-count = 6
hello-timer = 60
control-connect-establish-timer = 60
lac-incoming-call-timer = 60
base-udp-port = 0
dialout-auth-lns = no
dialout-send-profile-name = no
verify-remote-host-name = no
acct-tunnel-connection-encoding = normal
tunnel-server-pre-sccrq-lookup = no
```

| Parameter                       | Setting                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| first-retry-timer               | Initial interval, in milliseconds, that the system waits before retransmitting control packets in the attempt to establish an L2TP tunnel with an LNS. The valid range is from 100 to 5000. The default is 1000. For details about how this timer works with the <code>retry-count</code> parameter in establishing and maintaining tunnel sessions, see “Retry timers” on page 6-9.                                                                                                                                                                                                                                                                                                                                                                                                             |
| retry-count                     | Number of times the system retransmits control packets in the attempt to establish or reestablish a tunnel. The valid range is from 1 to 10. The default value is 10.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| hello-timer                     | Interval, in seconds, between Hello messages sent to the LNS. Specify a number from 0 to 600. The default is 60. The 0 setting specifies that no Hello messages are sent.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| control-connect-establish-timer | Number of seconds during which the system can establish an L2TP tunnel with another host. Enter an integer from 0 to 600. The default is 60.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| lac-incoming-call-timer         | Number of seconds the system waits for a session request to complete. Specify an integer from 1 to 600. The default is 60.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| base-udp-port                   | <p>UDP port number to use for the control link. The default zero value causes the system to dynamically assign a nonprivate port for exchanging control information while establishing a tunnel. This avoids the possibility of requesting a UDP port that is already in use. Nonzero values from 0 to 60,000 can be configured to use a known port, which is sometimes a firewall requirement. The configured value is used to assign a port number by using the following formula:</p> $\text{base-udp-port} + (\text{shelf-number} \times 100) + \text{slot-number}$ <p>A system reset is required for the L2TP subsystem to recognize the new UDP port number. The change does not take effect for existing active connections until they have been disconnected and then reestablished.</p> |
| dialout-auth-lns                | If set to <code>yes</code> , restricts the LAC to accept dialout requests only from the LNS which has authenticated during the tunnel setup.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| dialout-send-profile-name       | Enable/disable the LNS to send the CONNECTION profile name vendor-specific attribute along with the dialout request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

| Parameter                       | Setting                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| verify-remote-host-name         | Enable/disable verification of the host name returned by the LNS. When enabled, the LAC compares the hostname returned by the LNS in the Start-Control-Connection-Reply (SCCRP) packet to the server-auth-id value configured in the local TUNNEL-SERVER profile or the Tunnel-Server-Auth-ID value in a RADIUS profile. If the values do not match, the LAC terminates the tunnel request. |
| acct-tunnel-connection-encoding | Encoding method for the value of the RADIUS Acct-Tunnel-Connection attribute. The default normal setting specifies that the value consists of the source and destination IP addresses, tunnel ID, and connection ID. This value is expected by NavisRadius. See “CSN encoding for RADIUS accounting” on page 6-10.                                                                          |
| tunnel-server-pre-sccrp-lookup  | Enables/disables a lookup for a TUNNEL-SERVER profile when a password is not available, before initiating a tunnel request. The lookup must be enabled when using a TUNNEL-SERVER profile. See “Sample global tunneling configuration” on page 6-10.                                                                                                                                        |

## Retry timers

RFC 2661, *Layer Two Tunneling Protocol “L2TP”* specifies that each retransmission of an L2TP control message must employ an exponential backoff interval—if the first retransmission occurs after 1 second, the next retransmission should occur after 2 seconds, the next after 4 seconds, and so forth. The intention of this requirement is to restore normal traffic quickly after a control packet is lost. However, the exponential backoff algorithm can take many seconds to restore traffic if several control packets are dropped.

The Stinger system uses a more aggressive retransmission strategy than required by RFC 2661. Retransmission delay is calculated on the basis of the values specified in the L2-TUNNEL-GLOBAL profile of the first-retry-timer and retry-count parameters to provide a more linear backoff interval. This algorithm allows for a quicker recovery in the event of lost control packets. For example, with the following default values, retransmission occurs for up to 27 seconds:

```
[in L2-TUNNEL-GLOBAL:l2tp-config]
first-retry-timer = 1000
retry-count = 10
```

The initial transmission and 10 retransmissions (as specified by the retry count) occur at the following intervals:

- 1 1 second
- 2 1 second
- 3 1 second
- 4 1 second
- 5 2 seconds
- 6 2 seconds

- 7 3 seconds
- 8 3 seconds
- 9 4 seconds
- 10 4 seconds
- 11 5 seconds

With the following configured values, retransmission occurs for up to 11 seconds:

```
[in L2-TUNNEL-GLOBAL:l2tp-config]
first-retry-timer = 1000
retry-count = 6
```

The initial transmission and six retransmissions occur at the following intervals:

- 1 1 second
- 2 1 second
- 3 1 second
- 4 1 second
- 5 2 seconds
- 6 2 seconds
- 7 3 seconds

## CSN encoding for RADIUS accounting

With a nondefault value, the `acct-tunnel-connection-encoding` parameter specifies how the 32-bit L2TP call serial number (CSN) should be encoded as the value of the RADIUS `Acct-Tunnel-Connection` attribute. It can be encoded as a decimal or hexadecimal string.

For an L2TP tunnel itself, no CSN exists. With CSN encoding, the 32-bit value that is encoded for the tunnel itself represents the initiator tunnel ID in the low 16 bits and the server tunnel ID in the high 16 bits.



**Note** There is no guarantee that the L2TP call serial number (CSN) is unique at all times. Even in a single tunnel, identical CSN values can occur if tunnel links are initiated from both the LAC and LNS sides. Therefore, use caution when specifying CSN encoding in the `acct-tunnel-connection-encoding` parameter.

## Sample global tunneling configuration

The following commands enable the system to operate as LAC for L2TP tunnels. Tunnel authentication is enabled, and the name `stinger-lac` is used for authentication purposes if the client or TUNNEL-SERVER profile does not specify a name.

```
admin> read l2-tunnel-global
admin> set l2tp-mode = lac
admin> set l2tp-auth-enabled = yes
admin> set l2tp-system-name = stinger-lac
admin> set l2tp-config retry-count = 6
admin> write -f
```



This configuration limits the retransmission period for failed (or terminated) tunnel attempts to 11 seconds by reducing the retry count to 6. For details, see “Retry timers” on page 6-9.

## Configuring client connections with PPP authentication

If a PPP client’s profile is configured to initiate an L2TP tunnel, the LAC attempts to open a tunnel after it authenticates the client’s session request by means of a user name and password.

After the initial authentication, the LAC negotiates Link Control Protocol (LCP) with the client and then opens the PPP Auth state. The LAC then forwards relevant LCP information (*proxy LCP*) as well as the client’s name and password (*proxy authentication*) to the LNS, so the LNS does not need to restart negotiation.

With proxy LCP, the LAC sends the following information to the LNS:

- The first LCP Config Request packet received from the client
- The last LCP Config Request packet received from the client
- The last LCP Config Request packet the LAC sent to the client

With proxy authentication, the LAC initiates PPP authentication of the session request and then sends the client’s name and password to the LNS in the appropriate L2TP attribute-value pairs. The LNS can then complete PPP authentication.

## Overview of mobile client connection settings

Following are the parameters (shown with default values) for configuring a mobile client profile:

```
[in CONNECTION/""]
station* = ""
encapsulation-protocol = atm

[in CONNECTION/"" :pppoe-options]
pppoe = no

[in CONNECTION/pppoe:pppoe-options]
pppoe = yes

[in CONNECTION/"" :ppp-options]
recv-password = ""

[in CONNECTION/"" :tunnel-options]
profile-type = disabled
tunneling-protocol = l2tp-protocol
primary-tunnel-server = ""
secondary-tunnel-server = ""
```

| Parameter              | RADIUS attribute    | Setting                                                                                                                           |
|------------------------|---------------------|-----------------------------------------------------------------------------------------------------------------------------------|
| station                | User-Name (1)       | Name of the client device. The value is case sensitive, and must match the name the remote device presents during authentication. |
| encapsulation-protocol | Framed-Protocol (7) | Encapsulation protocol. Set to <b>ppp</b> for PPPoA or PPPoE.                                                                     |

## L2TP Tunneling Configuration

Configuring client connections with PPP authentication


---

| Parameter                                  | RADIUS attribute                            | Setting                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------------------------------|---------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| pppoe-options:<br>pppoe                    | Ascend-PPPoE-<br>Enable (74)                | Enables or disables processing of PPPoE packets. PPPoE is disabled by default. When both <code>pppoe</code> and <code>pppoa</code> are set to <code>yes</code> , the system automatically detects which protocol is used in the incoming traffic stream and performs the required processing. To enable only PPPoE processing, set this parameter to <code>yes</code> and the <code>pppoa</code> parameter to <code>no</code> . |
| pppoa-options:<br>pppoa                    | Ascend-PPPoA-<br>Enable (318)               | Enables or disables processing of PPPoA packets. PPPoA is enabled by default. When both <code>pppoae</code> and <code>pppoe</code> are set to <code>yes</code> , the system automatically detects which protocol is used in the incoming traffic stream and performs the required processing.                                                                                                                                   |
| ppp-options:<br>recv-password              | Password (2)                                | Password expected from the remote device to authenticate the session request.                                                                                                                                                                                                                                                                                                                                                   |
| tunnel-options:<br>profile-type            | NA                                          | Type of tunneling profile. Set to <code>mobile-client</code> for PPP clients using L2TP tunneling.                                                                                                                                                                                                                                                                                                                              |
| tunnel-options:<br>tunneling-protocol      | Tunnel-Type (64)<br>Tunnel-Medium-Type (65) | Protocol used to establish the tunnel. For L2TP tunneling, this setting must specify L2TP. The RADIUS value for <code>Tunnel-Medium-Type</code> must be IP (1).                                                                                                                                                                                                                                                                 |
| tunnel-options:<br>primary-tunnel-server   | Tunnel-Server-<br>Endpoint (67)             | IP address or hostname of a tunnel server. The system first attempts to establish a tunnel to the <code>primary-tunnel-server</code> . If it cannot reach the primary server, the system then attempts to establish a tunnel to the <code>secondary-tunnel-server</code> .<br><br>If a DNS lookup returns several IP addresses, the system attempts to establish a tunnel to each address in turn.                              |
| tunnel-options:<br>secondary-tunnel-server | Tunnel-Server-<br>Endpoint (67)             | IP address or hostname of a tunnel server. The system attempts to establish a tunnel to the <code>secondary-tunnel-server</code> only if its initial attempt to reach the <code>primary-tunnel-server</code> fails.<br><br>If a DNS lookup returns several IP addresses, the system attempts to establish a tunnel to each address in turn.                                                                                     |

## Sample PPP-authenticated client connection

In the following sample setup, the Stinger negotiates a PPPoA session from a DSL client, including user password authentication. It then initiates the L2TP tunnel, passing the authentication and LCP information on to an LNS across the Gigabit Ethernet interface.

Figure 6-2. Bringing up a tunnel using PPP authentication



The following commands disable tunnel authentication in the L2-TUNNEL-GLOBAL profile:

```
admin> read l2-tunnel-global
admin> set l2tp-mode = lac
admin> set l2tp-auth-enabled = no
admin> write -f
```

The following commands configure the IP-INTERFACE profile for the Gigabit Ethernet interface:

```
admin> read ip-interface { 1 8 2 }
admin> set ip-address = 1.1.1.200
admin> write -f
```

The next commands enable an SDSL interface in slot 4 and configure an ATM circuit from the interface to the LIM's ATM internal interface:

```
admin> read sdsl { 1 4 1 }
admin> set enabled = yes
admin> write -f
admin> new connection cir-141
admin> set active = yes
admin> set atm-options atm1483type = aal5-vc
admin> set atm-options vci = 36
admin> set atm-options nailed-group = 151
admin> set atm-connect-options nailed-group = 2097
admin> set atm-connect-options atm1483type = aal5-vc
admin> write -f
```

## **L2TP Tunneling Configuration**

*Configuring client connections with PPP authentication*

---

Following is a comparable RADIUS profile:

```
permconn-ran-1 Password = "ascend", Service-Type = Outbound-User,
 Framed-Protocol = ATM-CIR,
 User-Name = cir-141,
 Ascend-Route-IP = Route-IP-No,
 Ascend-ATM-Vci = 36,
 Ascend-ATM-Vpi = 0,
 Ascend-ATM-Group = 151,
 Ascend-ATM-Connect-Vci = 35,
 Ascend-ATM-Connect-Vpi = 0,
 Ascend-ATM-Connect-Group = 2097
```

The following commands create a CONNECTION profile for the mobile client in which the client's IP address (3.3.3.3) is assigned by the LNS when the session is established.

```
admin> new connection pppoa-1
admin> set active = yes
admin> set encapsulation-protocol = ppp
admin> set telco-options call-type = off
admin> set ppp-options recv-password = localpw
admin> set tunnel-options profile-type = mobile-client
admin> set tunnel-options primary-tunnel-server = 1.1.1.1
admin> set tunnel-options tunneling-protocol = l2tp
admin> write -f
```


Following is a comparable RADIUS profile:

```
pppoa-1 Password = "localpw"
 Service-Type = Framed-User,
 Framed-Protocol = PPP,
 Tunnel-Server-Endpoint = "1.1.1.1",
 Tunnel-Type = L2TP,
 Tunnel-Medium-Type = IP
```

## **Sample PPP-authenticated connection with two LNS systems**

Figure 6-3 shows a Stinger unit that can connect to one of two possible LNS end points to create an L2TP tunnel for the mobile client. In this example, the LNS end points are on remote networks, so the system requires a connection or RADIUS profile to establish a connection to one of the end-point systems.

Figure 6-3. Primary and secondary tunnel end points



The following commands configure the Stinger unit's system IP address:

```
admin> read ip-global
admin> set system-ip-addr = 2.2.2.2
admin> write -f
```

The following commands configure the IP-INTERFACE profile for the Gigabit Ethernet interface:

```
admin> read ip-interface { 1 8 2 }
admin> set ip-address = 1.1.1.200
admin> write -f
```

The next commands enable an SDSL interface in slot 4 and configure an ATM circuit from the interface to the LIM's ATM internal interface.

```
admin> read sdsl { 1 4 1 }
admin> set enabled = yes
admin> write -f
admin> new connection cir-141
admin> set active = yes
admin> set atm-options atm1483type = aal5-vc
admin> set atm-options vci = 36
admin> set atm-options nailed-group = 151
admin> set atm-connect-options nailed-group = 2097
admin> set atm-connect-options atm1483type = aal5-vc
admin> write -f
```

The following commands create a CONNECTION profile for the mobile client. In this case, the client's IP address is assigned in the profile, but that is not a requirement.

```
admin> new connection pppoa-1
admin> set active = yes
admin> set encapsulation-protocol = ppp
admin> set telco-options call-type = off
admin> set ppp-options recv-password = localpw
```

```

admin> set ip-options remote-addr = 3.3.3.3
admin> set tunnel-options profile-type = mobile-client
admin> set tunnel-options tunneling-protocol = l2tp
admin> set tunnel-options primary-tunnel-server = 1.1.1.1
admin> set tunnel-options secondary-tunnel-server = 4.4.4.4
admin> write -f

```

## **Configuring connection-based tunnel authentication**

Connection-based tunnel authentication uses values specified in the client's CONNECTION profile to authenticate a tunnel request. For information about how the system determines whether it must start up a new tunnel for the client or if it can use an existing tunnel to the destination network, see "Overview of L2TP tunnel authentication" on page 6-3.

### **Overview of connection-based tunnel authentication settings**

Following are the parameters (shown with default values) for configuring connection-based tunnel authentication. The client connection parameters defined in "Configuring client connections with PPP authentication" on page 6-11 also apply.

```

[in CONNECTION/"" :tunnel-options]
profile-type = disabled
tunneling-protocol = atmp-protocol
primary-tunnel-server = ""
secondary-tunnel-server = ""
password = ""
client-auth-id = ""
server-auth-id = ""
vrouter = ""
assignment-id = ""

```


| <b>Parameter</b>        | <b>RADIUS attribute</b>                     | <b>Setting</b>                                                                                                                                                                                                                                                          |
|-------------------------|---------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| profile-type            | NA                                          | Type of tunneling profile. Set to <code>mobile-client</code> for PPP clients using L2TP tunneling.                                                                                                                                                                      |
| tunneling-protocol      | Tunnel-Type (64)<br>Tunnel-Medium-Type (65) | Protocol used to establish the tunnel. For L2TP tunneling, this setting must specify L2TP. The RADIUS value for Tunnel-Medium-Type must be IP (1).                                                                                                                      |
| primary-tunnel-server   | Tunnel-Server-Endpoint (67)                 | IP address or hostname of a tunnel end point. If a DNS lookup returns several IP addresses, the system attempts to establish a tunnel to each address in turn. If the primary server is unavailable, the system attempts to establish a tunnel to the secondary server. |
| secondary-tunnel-server | Tunnel-Server-Endpoint (67)                 | IP address or hostname of a tunnel end point. If a DNS lookup returns several IP addresses, the system attempts to establish a tunnel to each address in turn. If the primary server is unavailable, the system attempts to establish a tunnel to the secondary server. |

| Parameter      | RADIUS attribute                | Setting                                                                                                                                                                                                                                                                                                   |
|----------------|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| password       | Tunnel-Password (69)            | Password used for authenticating the tunnel.                                                                                                                                                                                                                                                              |
| client-auth-id | Tunnel-Client-Auth-ID (90)      | Name sent to the tunnel server for authenticating the tunnel. The name can contain up to 31 characters. The Tunnel-Client-Auth-ID attribute can be specified in Access-Response packets and are also generated in Accounting-Request packets. See the <i>TAOS RADIUS Guide and Reference</i> for details. |
| server-auth-id | Tunnel-Server-Auth-ID (91)      | This field is currently ignored in a CONNECTION profile. The Tunnel-Server-Auth-ID attribute can be specified in Access-Response packets and are also generated in Accounting-Request packets. See the <i>TAOS RADIUS Guide and Reference</i> for details.                                                |
| vrouter        | Ascend-Tunnel-VRouter-Name (31) | Name of a virtual router to use for establishing the tunnel.                                                                                                                                                                                                                                              |
| assignment-id  | Tunnel-Assignment-ID (82)       | The tunnel assignment ID, as defined in RFC 2868, <i>RADIUS Attributes for Tunnel Protocol Support</i> . See “Using tunnel assignment IDs” on page 6-22.                                                                                                                                                  |

### Example of connection-based tunnel authentication

In the example shown in Figure 6-4, a mobile client dials into a Stinger unit to tunnel into its destination network across the Internet.

Figure 6-4. Connection-based tunnel authentication



The following commands enable tunnel authentication in the L2-TUNNEL-GLOBAL profile:

```
admin> read l2-tunnel-global
admin> set l2tp-mode = lac
admin> set l2tp-auth-enabled = yes
admin> write -f
```

The next commands enable an SDSL interface in slot 4 and configure an ATM circuit from the interface to the LIM's ATM internal interface.

```
admin> read sdsl { 1 4 1 }
admin> set enabled = yes
admin> write -f
```

## **L2TP Tunneling Configuration**

### *Configuring server-based tunnel authentication*

---

```
admin> new connection cir-141
admin> set active = yes
admin> set atm-options atm1483type = aal5-llc
admin> set atm-options vci = 36
admin> set atm-options nailed-group = 151
admin> set atm-connect-options nailed-group = 2097
admin> set atm-connect-options atm1483type = aal5-llc
admin> write -f
```

The following commands configure a local CONNECTION profile for the PPP client and specify a client-auth-id name:

```
admin> new connection ppp-user
admin> set encapsulation-protocol = ppp
admin> set ppp-options recv-password = localpw
admin> set telco-options call-type = off
admin> set tunnel-options profile-type = mobile-client
admin> set tunnel-options tunneling-protocol = l2tp-protocol
admin> set tunnel-options primary-tunnel-server = 1.1.1.1
admin> set tunnel-options password = lac-pass
admin> set tunnel-options client-auth-id = stinger-lac
admin> write -f
```

Note that you need not assign an IP address because it is assigned by the LNS. Following is a comparable RADIUS profile:

```
ppp-user User-Password = "localpw", Service-Type = Framed-User
 Tunnel-Type = L2TP,
 Tunnel-Server-Endpoint = 1.1.1.1,
 Tunnel-Password = lac-pass,
 Tunnel-Client-Auth-ID = stinger-lac
```

With the sample profiles, the LAC authenticates the PPP client's session request. It then initiates a tunnel to the LNS if a tunnel does not already exist to that end-point address. When the Stinger unit requests the tunnel, it sends the name `stinger-lac` and the password `lac-pass` to authenticate the tunnel. The LNS uses the same strings to authenticate the LAC before establishing the tunnel.

## **Configuring server-based tunnel authentication**

Server-based tunnel authentication uses values specified in a TUNNEL-SERVER profile to authenticate a tunnel request. For information about how the system determines whether it must start up a new tunnel for the client or if it can use an existing tunnel to the destination network, see "Overview of L2TP tunnel authentication" on page 6-3.



## Overview of server-based tunnel authentication settings

Following are the parameters (shown with default values) for configuring the Stinger unit to use server-based tunnel authentication:

```
[in TUNNEL-SERVER/""]
server-endpoint* = ""
enabled = yes
shared-secret = ""
client-auth-id = ""
server-auth-id = ""

[in L2-TUNNEL-GLOBAL]
l2tp-system-name = ""

[in L2-TUNNEL-GLOBAL:l2tp-config]
tunnel-server-pre-sccrq-lookup = no
```

| Parameter                                           | Setting                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| tunnel-server:<br>server-endpoint                   | Hostname or IP address that identifies the LNS. Usually, this is the same value as the Tunnel-Server-Endpoint RADIUS attribute, but it can differ. If a hostname is specified, the Stinger unit performs a DNS lookup for the host's address.                                                                                                                                                              |
| tunnel-server:<br>enabled                           | Enable/disable tunnels to the specified LNS.                                                                                                                                                                                                                                                                                                                                                               |
| tunnel-server:<br>shared-secret                     | Value (up to 21 characters) used by both the LAC and the LNS ends of the tunnel to authenticate tunnel requests initiated by local CONNECTION profiles. See "Example TUNNEL-SERVER password configuration" on page 6-20.                                                                                                                                                                                   |
| tunnel-server:<br>client-auth-id                    | LAC system name used for tunnel authentication. This name is sent in the Start-Control-Connection-Request (SCCRQ) message.                                                                                                                                                                                                                                                                                 |
| tunnel-server:<br>server-auth-id                    | LNS system name used for tunnel authentication. This name is sent in the SCCRP message.                                                                                                                                                                                                                                                                                                                    |
| l2-tunnel-global:<br>l2tp-system-name               | Name sent to the tunnel server for authenticating the tunnel if client-auth-id is not specified. See "System name used for tunnel authentication" on page 6-4.                                                                                                                                                                                                                                             |
| l2-tunnel-global:<br>tunnel-server-pre-sccrq-lookup | By default, the LAC initiates a tunnel request with the authentication parameters it has available. After it receives an L2TP SCCRP packet from the LNS, if further authentication is required, it looks for a TUNNEL-SERVER profile for the LNS.<br><br>If the tunnel-server-pre-sccrq-lookup parameter is set to yes, the LAC looks for a matching TUNNEL-SERVER profile before sending an SCCRQ packet. |

The Tunnel-Password (69) RADIUS attribute-value pair provides the same capability as the shared-secret parameter in a TUNNEL-SERVER profile. Its value is the password sent to the LNS for authenticating a tunnel request initiated from a

RADIUS user profile or a pseudo-user profile used for authentication only. The Tunnel-Password value must be encrypted by the RADIUS server. Otherwise, tunnel authentication fails.

### **Example TUNNEL-SERVER password configuration**

The following commands specify the password to be sent to the LNS for authenticating tunnels initiated from local CONNECTION profiles:

```
admin> read tunnel-server 1.1.1.1
admin> set enabled = yes
admin> set shared-secret = lac-pass
admin> write -f
```

Following is a RADIUS user profile that uses the Tunnel-Password attribute for authentication with the same LNS:

```
l2tp-client Password = "my-password"
 Service-Type = Framed-User,
 Framed-Protocol = PPP,
 Framed-IP-Address = 3.3.3.3,
 Framed-IP-Netmask = 255.255.255.255,
 Tunnel-Type = L2TP,
 Tunnel-Medium-Type = IP,
 Tunnel-Server-Endpoint = "1.1.1.1",
 Tunnel-Password = "lac-pass"
```

If you prefer, you can remove the Tunnel-Password attribute from clients' profiles and create a profile whose sole purpose is to authenticate L2TP tunnels. This approach causes an extra RADIUS lookup the first time the tunnel is created, but it simplifies administration when shared secrets change. The RADIUS profile for tunnel authentication must specify the L2TP peer's name, a null password (""), and the Outbound-User setting for Service-Type. For example:

```
1.1.1.1 Password = "", Service-Type = Outbound-User
 Tunnel-Password = "lac-pass"
```

When an L2TP tunnel is initially established, both the LNS and the LAC issue a RADIUS lookup based on the peer's name. If the system finds a profile such as the one shown in the preceding example, it uses the Tunnel-Password value to authenticate the tunnel.




**Note** The password in the pseudo-user profile must be null (""). Because the null password represents a security risk, the profile must specify the Outbound-User setting for Service-Type.

## Sample server-based authentication configuration

In the following example setup, the Stinger uses a TUNNEL-SERVER profile to obtain the tunnel authentication information for the LNS at 1.1.1.1.

Figure 6-5. Server-based tunnel authentication



The following commands enable tunnel authentication and enable a lookup for a TUNNEL-SERVER profile before initiating a tunnel request:

```
admin> read l2-tunnel-global
admin> set l2tp-mode = lac
admin> set l2tp-auth-enabled = yes
admin> set l2tp-system-name = stinger-lac
admin> set l2tp-config tunnel-server-pre-sccrq-lookup = yes
admin> write -f
```

The following commands configure a TUNNEL-SERVER profile for the LNS:

```
admin> new tunnel-server 1.1.1.1
admin> set enabled = yes
admin> set shared-secret = lac-pass
admin> set client-auth-id = stinger-lac
admin> write -f
```

Following is a comparable RADIUS profile:

```
1.1.1.1 Password = "", Service-Type = Framed-User
 Tunnel-Password = lac-pass,
 Tunnel-Client-Auth-ID = stinger-lac
```

The next commands enable an SDSL interface in slot 4 and configure an ATM circuit from the interface to the LIM's ATM internal interface.

```
admin> read sdsl { 1 4 1 }
admin> set enabled = yes
admin> write -f
admin> new connection cir-141
admin> set active = yes
admin> set atm-options atm1483type = aal5-llc
admin> set atm-options vci = 36
```

```
admin> set atm-options nailed-group = 151
admin> set atm-connect-options nailed-group = 2097
admin> set atm-connect-options atm1483type = aal5-llc
admin> write -f
```

The following commands configure a CONNECTION profile for the mobile client and do not specify a tunnel password or client-auth-id name:

```
admin> new connection ppp-user
admin> set active = yes
admin> set encapsulation-protocol = ppp
admin> set telco-options call-type = off
admin> set ppp-options recv-password = localpw
admin> set tunnel-options profile-type = mobile-client
admin> set tunnel-options tunneling-protocol = l2tp-protocol
admin> set tunnel-options primary-tunnel-server = 1.1.1.1
admin> write -f
```

Following is a comparable RADIUS profile:

```
ppp-user User-Password = "localpw", Service-Type = Framed-User
 Tunnel-Type = L2TP,
 Tunnel-Server-Endpoint = 1.1.1.1,
```

## ***Using tunnel assignment IDs***

Like the client-auth-id value, the tunnel assignment ID helps the LAC determine whether to assign a client session to an existing tunnel or to create a new one. This value is also used for grouping client sessions into specific tunnels. For details, see RFC 2868, *RADIUS Attributes for Tunnel Protocol Support*.

### ***Overview of tunnel assignment settings***

The following parameter (shown with default values) is used to specify a tunnel assignment ID:


```
[CONNECTION/"" :tunnel-options]
assignment-id = ""
```

| Parameter     | RADIUS attribute          | Setting                                                                                                                                                                                                  |
|---------------|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| assignment-id | Tunnel-Assignment-ID (82) | Identification (name) assigned to tunnels to allow grouping sessions. A text string of up to 31 characters. The value has local significance only. It is not transmitted to the remote tunnel end point. |

## Example of configuring tunnel assignment IDs

In this example, the Stinger unit is configured to perform tunnel authentication for L2TP tunnels. The two PPP clients shown in Figure 6-6 are configured to use different tunnels to the LNS on the basis of their tunnel assignment IDs. (The same clients can be configured to use the same multiplexed tunnel if their tunnel assignment IDs are set to the same string.)

Figure 6-6. Tunnel assignment IDs



The following set of commands enables tunnel authentication:

```
admin> read l2-tunnel-global
admin> set l2tp-mode = lac
admin> set l2tp-auth-enabled = yes
admin> write -f
```

The next commands enable two SDSL interfaces in slot 4 and configure an ATM circuit from each interface to the LIM's ATM internal interface:

```
admin> read sdsl { 1 4 1 }
admin> set enabled = yes
admin> write -f
admin> read sdsl { 1 4 2 }
admin> set enabled = yes
admin> write -f
admin> new connection cir-141
admin> set active = yes
admin> set atm-options atm1483type = aal5-llc
admin> set atm-options vci = 36
admin> set atm-options nailed-group = 151
admin> set atm-connect-options nailed-group = 2097
admin> set atm-connect-options atm1483type = aal5-llc
admin> write -f
```

## **L2TP Tunneling Configuration**

*Using tunnel assignment IDs*

---

```
admin> new connection cir-142
admin> set active = yes
admin> set atm-options atm1483type = aal5-llc
admin> set atm-options vci = 36
admin> set atm-options nailed-group = 152
admin> set atm-connect-options nailed-group = 2097
admin> set atm-connect-options atm1483type = aal5-llc
admin> write -f
```

The following set of commands creates CONNECTION profiles for the two mobile clients:

```
admin> new connection sdsuser-1
admin> set encapsulation-protocol = ppp
admin> set ppp-options recv-password = localpw
admin> set telco-options call-type = off
admin> set tunnel-options profile-type = mobile-client
admin> set tunnel-options tunneling-protocol = l2tp-protocol
admin> set tunnel-options primary-tunnel-server = 1.1.1.1
admin> set tunnel-options password = lac-pass
admin> set tunnel-options client-auth-id = stinger-lac
admin> set tunnel-options assignment-id = sds11-taid
admin> write -f

admin> new connection sdsuser-2
admin> set encapsulation-protocol = ppp
admin> set ppp-options recv-password = localpw
admin> set telco-options call-type = off
admin> set tunnel-options profile-type = mobile-client
admin> set tunnel-options tunneling-protocol = l2tp-protocol
admin> set tunnel-options primary-tunnel-server = 1.1.1.1
admin> set tunnel-options password = lac-pass
admin> set tunnel-options client-auth-id = stinger-lac
admin> set tunnel-options assignment-id = sds12-taid
admin> write -f
```

Following are comparable RADIUS profiles:

```
sdsuser-1 Password = "localpw"
 User-Service = Framed-User,
 Framed-Protocol = PPP,
 Test-Idle-Limit = 0,
 Tunnel-Type = L2TP :1,
 Tunnel-Server-Endpoint = 1.1.1.1 :1,
 Tunnel-Client-Auth-ID = stinger-lac: 1,
 Tunnel-Password = lac-pass,
```

```
Tunnel-Assignment-ID = sds11-taid:1
sds1user-2 Password = "localpw"
User-Service = Framed-User,
Framed-Protocol = PPP,
Test-Idle-Limit = 0,
Tunnel-Type = L2TP :1,
Tunnel-Server-Endpoint = 1.1.1.1 :1,
Tunnel-Client-Auth-ID = stinger-lac: 1,
Tunnel-Password = lac-pass,
Tunnel-Assignment-ID = sds12-taid:1
```

RADIUS accounting Stop records display the Tunnel-Assignment-ID used for the user session. For example:

```
Tue May 2 15:58:08 2003
User-Name = "sds1user-1"
NAS-Identifier = 2.2.2.2
NAS-Port = 11313
NAS-Port-Type = Async
Acct-Status-Type = Stop
Acct-Delay-Time = 0
Acct-Session-Id = "317658341"
Acct-Authentic = Local
Acct-Session-Time = 112
Acct-Input-Octets = 2155
Acct-Output-Octets = 513
Acct-Input-Packets = 23
Acct-Output-Packets = 14
Ascend-Disconnect-Cause = 185
Ascend-Connect-Progress = 60
Ascend-Xmit-Rate = 28800
Ascend-Data-Rate = 33600
Ascend-PreSession-Time = 19
Ascend-Pre-Input-Octets = 0
Ascend-Pre-Output-Octets = 0
Ascend-Pre-Input-Packets = 0
Ascend-Pre-Output-Packets = 0
Ascend-Modem-PortNo = 1
Ascend-Modem-SlotNo = 7
Ascend-Modem-ShelfNo = 1
Tunnel-Type = L2TP
Tunnel-Server-Endpoint = "1.1.1.1"
Tunnel-Client-Auth-ID = "stinger-lc"
Tunnel-Server-Auth-ID = "dest-lns"
Tunnel-Assignment-ID = "sds11-taid"
```

## Administrative tools for monitoring L2TP

The stats command, described in detail in the *Stinger Reference*, now provides options for gathering and displaying traffic statistics related to L2TP and PPPoE connections. Following is the syntax of the stats command in the context of L2TP monitoring:

```
stats stats-cmd np stats-type
```

| Command element | Description                                                                                                                                                                                                                                                                                               |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| stats-cmd       | One of the following actions: <ul style="list-style-type: none"><li>■ traffic (Display traffic statistics.)</li><li>■ clear (Clear statistics by zeroing out the buffer.)</li></ul>                                                                                                                       |
| np stats-type   | Type of information to display about network processor operations. Specify one of the following options: <ul style="list-style-type: none"><li>■ l2tp (Display L2TP information.)</li><li>■ pppoe (Display PPPoE information.)</li><li>■ ipport (Display information about open UDP/TCP ports.)</li></ul> |

### Displaying traffic statistics for open UDP/TCP ports

For example, the following command displays traffic statistics for open UDP/TCP ports:

```
admin> stats traffic np ipport
IP Port NP statistics
```

```
~~~~~
Rx UDP packets for us:                328
Rx rogue UDP (no listening port):      1
```

The following command clears ipport statistics:

```
admin> stats clear np ipport
```

### Displaying L2TP statistics

The following command displays network processor L2TP statistics:

```
admin> stats traffic np l2tp
L2TP NP statistics
```

```
~~~~~
Upstream (MC -> LNS):

PPPoA: 0
PPPoE: 288
tunnel DID not found: 0
last tunnel DID matched: 505
*** fast path (Tx RSP) ***: 288
```



Downstream (LNS -> MC):

```

Rx L2TP Packets: 336 [336 ETH, 0 ATM]
 control: 18
 data: 318
 - with Length: 0
 - with Sequencing: 0
 - with Offset: 0
 - bad Version: 0
offset overflow: 0
mobile DID not found: 10
last mobile DID matched: 504
*** fast path (Tx RSP) ***: 308
```

## **Displaying PPPoE statistics**

The following command displays statistics about PPPoE sessions on the network processor:

```
admin> stats traffic np pppoe
PPPoE NP parsing statistics
```

```
~~~~~  
Rx PPPoE packets:                439  
Incorrect LLC header:            0  
Discovery packet (slow path):    4  
Incorrect ETHER-TYPE:            0  
Incorrect VERSION/TYPE:          0  
Unknown Session:                 20  
Locally terminated:              120  
Not PPP Session stage (slow path): 0
```

## **Limitations with this software version**

The following limitations apply to the Stinger L2TP feature set:

- Using L2TP sequence numbers on the data channel is not supported.
- To take advantage of the IP control module wire-speed capability, the LNS must be reachable through one of the following interfaces:
  - Gigabit Ethernet interface
  - Gigabit Ethernet VLAN interface
  - IP over ATM connection on a trunk interface using AAL5 LLC
- Use of the 48-port ADSL LIMs Annex A (STGR-LIM-AD-48 or STGRRT-LIM-AD-48), Annex B (STGR-LIM-AB-48), and Annex C (STGR-LIM-AC-48) is not supported for L2TP.
- IP QoS mechanisms do not currently apply to L2TP calls.

### ***Impact of L2TP data sequencing limitation***

The Stinger acting as LAC will never insert sequence numbers in data messages. If it receives an L2TP data message with sequence numbers in it, the system processes the data stream but will not begin sending sequence numbers in future outgoing data messages. This behavior is not compliant with RFC 2661, section 5.4.

The impact of not sending sequence numbers in data packets depends on the LNS vendor implementation. Some LNS systems accept this behavior, and others may begin to drop packets. Some LNS systems drop only a fraction of the packets, so troubleshooting the problem can be difficult. For this reason, Alcatel-Lucent strongly recommends that you turn off sequencing of the L2TP data channels on the LNS when interacting with the Stinger as LAC.

### ***Impact of L2TP IP QoS limitation***

IP QoS (rate limiting) is currently not supported for L2TP connections. This means that if a PACKET-FLOWS profile is applied to an L2TP CONNECTION profile or even to an egress interface, the rules are ignored for L2TP connections. This applies to both upstream and downstream traffic.

---

# Virtual Router Configuration

# 7

|                                                    |      |
|----------------------------------------------------|------|
| Overview of virtual routing . . . . .              | 7-1  |
| Creating a virtual router. . . . .                 | 7-3  |
| Deleting a virtual router. . . . .                 | 7-12 |
| Administrative tools for virtual routers . . . . . | 7-12 |


Virtual routing enables you to partition virtual domains within the global IP router. Each virtual domain is defined by a named virtual router. Currently, a Stinger IP DSLAM supports up to 1022 virtual routers in addition to the global router.

If you do not configure virtual routers, the global IP router operates as documented in Chapter 4, “IP Routing Configuration.” All interfaces that are not explicitly grouped with a defined virtual router are grouped with the global router.

## Overview of virtual routing

Virtual routing enables high-density circuit termination with secure logical partitioning and multiple route tables. Virtual routing is particularly useful for remote access server (RAS) functionality. For example, it can securely partition traffic from many CPE devices to different ISPs, with each ISP mapped to a separate virtual domain. Figure 7-1 shows a simplified diagram with three virtual routers configured in the global router. Because each virtual router has its own routing table, traffic within a virtual domain is typically not shared with other domains.

Figure 7-1. Simple diagram of three virtual domains (virtual routers)



## **How virtual routers affect the routing table**

When virtual routers are not defined, the global router maintains a single IP routing table that enables the router to reach any of its many interfaces. In that context, each interface known to the system requires a unique address.

With virtual routers, addresses must be unique within the virtual domain, but not necessarily within the Stinger system. Because each virtual router maintains its own routing table, and because it knows about only those interfaces that explicitly specify the same virtual router, there is no requirement that the private networks maintain unique address spaces, as long as the virtual domains are not interconnected.

## **Interconnecting virtual domains**

Each virtual router has its own associated routing table, ARP table, route cache, and address pools, which cannot be shared with another virtual router. However, it is possible to interconnect two virtual domains by defining an inter-virtual-router route. For details, see “Specifying an inter-virtual-router route” on page 7-10.



**Note** Since routing traffic between virtual domains is not fast routed, it is strongly recommended to minimize such traffic.

## **Applicability and limitations**

When configuring virtual routing on a Stinger IP DSLAM, consider the following issues related to virtual router applicability and limitations:

- Virtual routing does not apply to switched or bridged connections.
- Only terminated virtual circuits (IP, PPPoA, PPPoE) can be integrated into a virtual domain.
- When the virtual LAN (VLAN) feature is used in its usual bridging capacity, to form a bridge between a VLAN ID on Gigabit Ethernet and a DSL interface, virtual routing does not apply. If the VLAN is not used in a bridge capacity but acts as an IP interface, virtual routing applies.
- Bridged IP routing (BIR) profiles have no special considerations for virtual routing. The CONNECTION profile terminating the PVC must define the virtual router.
- With the current implementation, multicasting applications must use the global router.
- If multiple virtual IP interfaces are configured on a physical interface, particularly an Ethernet interface, all virtual interfaces must be attached to the same virtual router.
- SNMP management utilities do not currently display information on a per-virtual-router basis.
- Errors and events are not logged on a per-virtual-router basis.
- The syslog host defined in the system’s LOG profile must be accessible to the global router.
- Servers defined in the DEBUG, TRAP, EXTERNAL-AUTH, IP-GLOBAL (for SNTp and multicast), CALL-LOGGING, and SNMP profiles must be accessible to the global router.
- Rip updates are not sent on all routed-vlans configured as vrouters.

## Creating a virtual router

When at least one VROUTER profile is configured, the `system-ip-address` parameter and the `global-vrouter` parameter in the IP-GLOBAL profile apply to the global router. All interfaces that are not explicitly assigned to another virtual router are grouped with the global router.

For each virtual router in the system, an instance of RIP is created to process routes. The new instance of RIP sends and receives update packets only on the interfaces associated with its particular virtual router and manipulates only that virtual router's routing table. A default instance of RIP is always created for the global router.

When you create a virtual router, the new instance of RIP sends and receives packets only on the interfaces associated with that virtual router and manipulates only that virtual router's routing table. All RIP-related parameters in a VROUTER profile use default settings that are recommended for most sites.

## Overview of VROUTER profile settings

A VROUTER profile contains the following parameters, shown with default values:

```
[in VROUTER/""]  
name* = ""  
active = yes  
vrouter-ip-addr = 0.0.0.0  
pool-base-address = [ 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0+  
assign-count = [ 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 +  
pool-name = [ "" "" "" "" "" "" "" "" "" "" "" "" "" "" "" "" "" "" "" "" +  
pool-summary = no  
share-global-pool = yes  
rip-policy = Poison-Rvrs  
summarize-rip-routes = no  
rip-trigger = yes  
domain-name = ""  
sec-domain-name = ""  
dns-primary-server = 0.0.0.0  
dns-secondary-server = 0.0.0.0  
client-primary-dns-server = 0.0.0.0  
client-secondary-dns-server = 0.0.0.0  
allow-as-client-dns-info = True
```

| Parameter          | Setting                                                                                                                                                                                  |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| name               | Unique name for the virtual router, up to 15 characters.<br>All interfaces belonging to a virtual router specify the same virtual router name in the IP-INTERFACE or CONNECTION profile. |
| active             | Activate the virtual router.                                                                                                                                                             |
| vrouter-ip-address | System IP address for the virtual router.                                                                                                                                                |

| Parameter            | Setting                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| pool-base-address    | Base address of a pool of contiguous addresses on a local network or subnet. The pool will be exclusively for use by the virtual router. For details about defining address pools, see “Configuring and using address pools” on page 4-23. For information about associating an address pool with a slot and virtual router combination, see “Slot-based address assignment” on page 4-36. |
| assign-count         | Number of addresses in the pool. The pool will be exclusively for use by the virtual router. For details about defining address pools, see “Configuring and using address pools” on page 4-23.                                                                                                                                                                                             |
| pool-name            | A pool name. Not currently used.                                                                                                                                                                                                                                                                                                                                                           |
| pool-summary         | Set/clear the pool summary flag to specify that the address pools will be summarized. For details about defining address pools, see “Configuring and using address pools” on page 4-23.                                                                                                                                                                                                    |
| share-global-pool    | Enable/disable the virtual router to share the address pools defined in the IP-GLOBAL profile.                                                                                                                                                                                                                                                                                             |
| rip-policy           | Policy for the virtual router to use when sending update packets that include routes received on the same interface. For details, see “Setting RIP options” on page 4-11.                                                                                                                                                                                                                  |
| summarize-rip-routes | Whether the virtual router summarizes subnet information in RIP-v1 advertisements. For details about this feature, see “Setting RIP options” on page 4-11.                                                                                                                                                                                                                                 |
| rip-trigger          | Enable/disable RIP triggering for the virtual router. For details about RIP triggering, see “Setting RIP options” on page 4-11.                                                                                                                                                                                                                                                            |



**Note** For details about `domain-name` and other DNS parameters, see “Configuring virtual router DNS servers” on page 7-10.

### **Example of defining a virtual router**

The following commands create a virtual router named `vr1` with a system address of `130.200.200.100`:

```
admin> new vrouter vr1
admin> set vrouter-ip-addr = 130.200.200.100
admin> write -f
```

The vr1 virtual router maintains minimal routing and interface tables at this point, as shown in the following sample output:

```
admin> netstat vr1 -rn
```

| Destination  | Gateway | IF      | Flg | Pref | Met | Use | Age  |
|--------------|---------|---------|-----|------|-----|-----|------|
| 127.0.0.0/8  | -       | bh0_vr1 | CP  | 0    | 0   | 1   | 6815 |
| 127.0.0.1/32 | -       | local   | CP  | 0    | 0   | 1   | 6815 |
| 127.0.0.2/32 | -       | rj0_vr1 | CP  | 0    | 0   | 1   | 6815 |
| 224.0.0.9/32 | -       | local   | CP  | 0    | 0   | 1   | 6815 |

Total Routes = 4      Hidden Routes = 0

```
admin> netstat vr1 -in
```

| Name    | MTU   | Net/Dest     | Address   | Ipkts | Ierr | Opkts | Oerr |
|---------|-------|--------------|-----------|-------|------|-------|------|
| vr0_vr1 | 1500  | 127.0.0.4/32 | 127.0.0.4 | 0     | 0    | 0     | 0    |
| lo0_vr1 | 1500  | 127.0.0.1/32 | 127.0.0.1 | 0     | 0    | 0     | 0    |
| local   | 65535 | 127.0.0.1/32 | 127.0.0.1 | 0     | 0    | 0     | 0    |
| rj0_vr1 | 1500  | 127.0.0.2/32 | 127.0.0.2 | 0     | 0    | 0     | 0    |
| bh0_vr1 | 1500  | 127.0.0.3/32 | 127.0.0.3 | 0     | 0    | 0     | 0    |

The virtual router also maintains its own IP, TCP, UDP, and ICMP statistics. For example:

```
admin> netstat vr1 -s
```

udp:

1442 packets received  
0 packets received with no ports  
0 packets received with errors  
0 packets dropped  
32 packets transmitted

tcp:

0 active opens  
1 passive opens  
0 connect attempts failed  
0 connections were reset  
1 connections currently established  
858 segments received  
0 segments received out of order  
548 segments transmitted  
0 segments retransmitted  
0 active closes  
0 passive closes  
0 disconnects while awaiting retransmission

icmp:

31 packets received  
0 packets received with errors  
Input histogram:  
30 echo requests  
1 netmask requests

31 packets transmitted

0 packets not transmitted due to lack of resources

Output histogram:

30 echo replies

1 netmask replies

ip:

0 packets received

0 packets received with header errors

0 packets received with address errors

0 packets received forwarded

0 packets received with unknown protocols

0 inbound packets discarded

0 packets delivered to upper layers

0 transmit requests

0 discarded transmit packets

0 outbound packets with no route

0 reassemblies timeout

0 reassemblies required

0 reassemblies succeeded

0 reassemblies failed

0 fragmentation succeeded

0 fragmentation failed

0 fragmented packets created

0 route discards due to lack of memory

64 default ttl

igmp:

0 packets received

0 bad checksum packets received

0 bad version packets received

0 query packets received

0 leave packets received

0 packets transmitted

0 query packets sent

0 response packets sent

0 leave packets sent

mcast:

0 packets received

0 packets forwarded

0 packets in error

0 packets dropped

0 packets transmitted

pim:

0 packets received

559 packets transmitted

559 hello packets sent



**Note** Multicast is not currently supported on a per-virtual-router basis, so the IGMP, multicast, and PIM statistics relate only to the global router.



## Defining address pools for a virtual router

For information about associating an address pool with a slot and virtual router combination, see “Slot-based address assignment” on page 4-36.

The following commands define an address pool for the `vr1` virtual router defined in “Example of defining a virtual router” on page 7-4:

```
admin> read vrouter vr1
admin> set pool-base 1 = 130.100.100.128
admin> set assign-count 1 = 127
admin> write -f
```

Following is a comparable RADIUS pool definition:

```
pools-taos01 Password = "ascend", Service-Type = Outbound-User
Ascend-IP-Pool-Definition = "1 130.100.100.128 127 vr1"
```

The `vr1` virtual router is now maintaining the following pool of addresses:

```
admin> ip-pools vr1
Pool#           Base           Count      InUse
  1           130.100.100.128       127         0
Number of remaining allocated addresses:  0
```



**Note** The `Ascend-IP-Pool-Definition` attribute supports a virtual router name as the last syntax element in a pool definition. The value of `Ascend-IP-Pool-Definition` uses the following syntax:

```
"pool-num base-addr assign-count [vrouter-name]"
```

For background information about address pools, see “Configuring and using address pools” on page 4-23. The process of defining address pools for a virtual router is the same as described in that section.

## Assigning interfaces to a virtual router

To assign virtual router membership to an interface, you specify a virtual router name in the interface profile. For a virtual router to be active, at least one IP interface (LAN or WAN) must specify its name.

### Overview of interface virtual router settings

To assign virtual router membership to an interface in local profiles, set the `vrouter` parameter. For example:

```
[in IP-INTERFACE/{ { shelf-1 slot-8 2 } 0 } ]
vrouter = vr1

[in CONNECTION/vr1-client]
vrouter = vr1
```

| Parameter            | RADIUS attribute                       | Setting                                                                                                                                                                |
|----------------------|----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>vrouter</code> | <code>Ascend-VRouter-Name (102)</code> | Name of a defined virtual router. Specifying the virtual router name groups the interface with the virtual router. The default null value specifies the global router. |

## Examples of assigning virtual router membership to interfaces

The following commands assigns a WAN interface to the `vr1` virtual router:

```
admin> read connection router-1
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set vrrouter = vr1
admin> set ip-options remote-address = 10.7.8.200/30
admin> set atm-options vci = 100
admin> set atm-options nailed-group = 201
admin> write -f
```

Following is a comparable RADIUS profile:

```
permconn-st-1 Password = "ascend"
    Service-Type = Outbound-User,
    Framed-Protocol = ATM-1483,
    User-Name = "router-1",
    Framed-IP-Address = 10.7.8.200,
    Framed-IP-Netmask = 255.255.255.252,
    Ascend-ATM-Group = 201,
    Ascend-Route-IP = Route-IP-Yes,
    Ascend-ATM-Vpi = 8,
    Ascend-ATM-Vci = 100,
    Ascend-Vrrouter-Name = "vr1"
```

For an example of assigning virtual router membership to a VLAN interface, see “Assigning a virtual router to a routed VLAN” on page 3-38.

## Defining virtual router static routes

You specify a static route associated with a virtual router for one of the following reasons:

- To define a route on a per-virtual-router basis
- To specify an inter-virtual-router route

## Overview of static route settings

Following are the virtual router static route parameters (shown here with default values) in IP-ROUTE profiles:

```
[in IP-ROUTE/""]  
vrrouter = ""  
inter-vrrouter = ""
```

| Parameter | Setting                                                                                                                                                                                           |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| vrrouter  | Name of the virtual router that will own this route. The route will be part of the specified virtual router's routing table. If no name is specified (the default), the global router is assumed. |

| Parameter     | Setting                                                                                                                                                                                                                                                                           |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| inter-vrouter | Name of a virtual router to use as the route's next hop. All packets to the static route's destination network are sent to the specified virtual router for a routing decision. The <b>gateway-address</b> parameter must be set to the zero address for this parameter to apply. |

In a RADIUS profile, the value of the Framed-Route (22) attribute can specify a virtual router name in the following syntax:

"dest-addr [/prefix] gateway-addr metric [private] [profile] [preference] [vrouter-name]"



**Note** The fields within the value of the Framed-Route attribute are positional. With the exception of the optional prefix-length specification, if any of the optional fields are specified, the optional fields to the left of that setting must also be specified.

## Examples of defining a route on a per-virtual-router basis

When you define a route on a per-virtual-router basis, it appears only in the specified virtual router's routing tables. That virtual router "owns" the route.

Following is an example of defining a static route within the **vr1** virtual router domain. This route will appear only in the routing table for **vr1**.

```
admin> new ip-route rtr1
admin> set dest = 10.5.6.7/28
admin> set gateway = 10.1.1.1
admin> set vrouter = vr1
admin> write -f
```

Following is a comparable RADIUS profile:

```
route-taos-1 Password = "ascend", Service-Type = Outbound-User
  Framed-Route = "10.5.6.7/28 10.1.1.1 7 n rtr1 60 vr1"
```

The following sample output shows the new static route that was added to the **vr1** virtual router's routing table:

```
admin> netstat vr1 -rn
```

| Destination  | Gateway  | IF      | Flg | Pref | Met | Use | Age  |
|--------------|----------|---------|-----|------|-----|-----|------|
| 10.1.1.0/24  | 10.1.1.1 | wan30   | SG  | 120  | 7   | 0   | 9    |
| 10.1.1.1/32  | 10.1.1.1 | wan30   | S   | 120  | 7   | 2   | 9    |
| 10.5.6.0/28  | 10.1.1.1 | wan30   | SG  | 60   | 8   | 0   | 9    |
| 11.1.1.0/24  | 11.1.1.1 | wan31   | SG  | 120  | 7   | 0   | 9    |
| 11.1.1.1/32  | 11.1.1.1 | wan31   | S   | 120  | 7   | 1   | 9    |
| 12.1.1.0/24  | 12.1.1.1 | wan32   | SG  | 120  | 7   | 0   | 9    |
| 12.1.1.1/32  | 12.1.1.1 | wan32   | S   | 120  | 7   | 1   | 9    |
| 127.0.0.0/8  | -        | bh0_vr1 | CP  | 0    | 0   | 0   | 2274 |
| 127.0.0.1/32 | -        | local   | CP  | 0    | 0   | 0   | 2274 |
| 127.0.0.2/32 | -        | rj0_vr1 | CP  | 0    | 0   | 0   | 2274 |

## Specifying an inter-virtual-router route

You can cause one virtual router or the global router to forward traffic to another virtual router for a routing decision by specifying an inter-virtual-router static route.



**Note** Because routing traffic between virtual domains is not fast routed, it is strongly recommended to minimize such traffic.

In the following example, the static route specifies the `vr1` virtual router as the route's next hop. This route is not defined on a per-virtual-router basis, so it is owned by the global router.

```
admin> new ip-route rtr2
admin> set dest-address = 11.0.0.0/24
admin> set inter-vrouter = vr1
admin> write -f
```

Following is a comparable RADIUS route profile:

```
route-taos-1 Password = "ascend", Service-Type = Outbound-User
    Framed-Route = "11.0.0.0/24 0.0.0.0 vr1"
```

The following output shows that the route has been added to the global router's routing table:

```
admin> netstat -rn
```

| Destination        | Gateway     | IF       | Flg | Pref | Met | Use  | Age  |
|--------------------|-------------|----------|-----|------|-----|------|------|
| 0.0.0.0/0          | 10.1.6.1    | ie0      | SGP | 60   | 1   | 59   | 4    |
| 11.0.0.0/24        | -           | vr0_vr1  | S   | 60   | 8   | 0    | 4    |
| 20.0.0.0/8         | -           | ie1-12-1 | C   | 0    | 0   | 12   | 234  |
| 20.1.1.2/32        | -           | local    | CP  | 0    | 0   | 0    | 2347 |
| 127.0.0.0/8        | -           | bh0      | CP  | 0    | 0   | 0    | 2378 |
| 127.0.0.1/32       | -           | local    | CP  | 0    | 0   | 0    | 2378 |
| 127.0.0.2/32       | -           | rj0      | CP  | 0    | 0   | 0    | 2378 |
| 130.1.1.1/32       | -           | sip0     | C   | 0    | 0   | 0    | 2378 |
| 130.1.1.252/30     | -           | rj0      | C   | 0    | 0   | 0    | 2378 |
| 100.1.6.0/24       | 100.1.6.221 | wanabe   | SG  | 60   | 1   | 0    | 4    |
| 101.1.6.0/24       | -           | ie0      | C   | 0    | 0   | 2531 | 2378 |
| 101.1.6.234/32     | -           | local    | CP  | 0    | 0   | 4152 | 2378 |
| 224.0.0.0/4        | -           | mcast    | CP  | 0    | 0   | 0    | 2378 |
| 224.0.0.1/32       | -           | local    | CP  | 0    | 0   | 0    | 2378 |
| 224.0.0.2/32       | -           | local    | CP  | 0    | 0   | 0    | 2378 |
| 224.0.0.5/32       | -           | local    | CP  | 0    | 0   | 732  | 2378 |
| 224.0.0.6/32       | -           | local    | CP  | 0    | 0   | 0    | 2378 |
| 255.255.255.255/32 | -           | ie0      | P   | 0    | 0   | 422  | 2378 |

## Configuring virtual router DNS servers

Virtual router DNS configuration includes settings for primary and secondary DNS servers, domain names, and client DNS servers. The settings direct connections that belong to the virtual router to a particular DNS service. To completely segment the virtual router's DNS information from any other hosts, you can configure and manage DNS information separately for each virtual router. The addresses configured for client DNS servers are presented to dial-in users during IP Control Protocol (IPCP) negotiation.

If DNS information is not found in the VROUTER profile, the system uses the DNS information in the IP-GLOBAL profile. The DNS list and the local DNS table maintained in RAM are systemwide DNS configurations that are not supported separately for each virtual router.

## Overview of virtual router DNS settings

Following are the virtual router-specific DNS parameters (shown with their default settings):

```
[in VROUTER/""]  
domain-name = ""  
sec-domain-name = ""  
dns-primary-server = 0.0.0.0  
dns-secondary-server = 0.0.0.0  
client-primary-dns-server = 0.0.0.0  
client-secondary-dns-server = 0.0.0.0  
allow-as-client-dns-info = True
```

| Parameter                   | Setting                                                                                                                                                             |
|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| domain-name                 | Primary domain name (up to 63 characters) to use for DNS lookups for this virtual router. The system appends this domain name to hostnames when performing lookups. |
| sec-domain-name             | Secondary domain name to use for DNS lookups for this virtual router if the hostname is not found in the primary domain.                                            |
| dns-primary-server          | Address of the primary local DNS server to use for lookups for this virtual router.                                                                                 |
| dns-secondary-server        | Address of the secondary local DNS server to use for lookups for this virtual router. Used only if the primary server is not found.                                 |
| client-dns-primary-server   | Address of a client DNS server for dial-in clients of this virtual router.                                                                                          |
| client-dns-secondary-server | Address of a secondary DNS server for dial-in clients of this virtual router.                                                                                       |
| allow-as-client-dns-info    | Enable/disable use of local DNS information if the client DNS servers are not found. To isolate local network information for this virtual router, set to false.    |

## Example of a typical virtual router DNS configuration

The following commands specify a primary and secondary domain name for DNS lookups for a virtual router named xyz:

```
admin> read vrouter xyz  
admin> set domain-name = xyz.com  
admin> set sec-domain-name = eng.xyz.com
```

```
admin> write -f
```

If a lookup fails in the first domain, the router tries again with the secondary domain name. To enable the system to use DNS to look up addresses, specify DNS server addresses, as shown in the following example:

```
admin> read vrouter xyz
admin> set dns-primary-server = 1.2.2.2
admin> set dns-secondary-server = 1.3.3.3
admin> write -f
```

If the primary server is unavailable, the system attempts a lookup on the secondary server. The following commands configure a client DNS server for this virtual router:

```
admin> read vrouter xyz
admin> set client-dns-primary-server = 1.2.2.2
admin> set client-dns-secondary-server = 1.2.2.96
admin> set allow-as-client-dns-info = false
admin> write -f
```

The secondary server is accessed only if the primary one is inaccessible. If both of these client DNS servers are not accessible, the system does not allow the client to access local DNS servers.

## ***Deleting a virtual router***

You can delete a virtual router only if no more interfaces are attached to it. If one interface is attached to a virtual router, the system prevents its deletion. To delete a virtual router that has no attached interfaces, delete the VROUTER profile. For example:

```
admin> delete vrouter vr1
```

Alcatel-Lucent recommends that you reset the system after deleting a virtual router with active connections. If a system reset is not possible, the recommended course of action before deleting the virtual router is to manually tear down its active connections, and then modify the local CONNECTION, IP-INTERFACE, and IP-ROUTE profiles that point to the virtual router to point instead to the global router or another existing virtual router.

## ***Administrative tools for virtual routers***

You can specify a virtual router name on the command line of the network administration commands listed in Table 7-1 to obtain information specific to a particular virtual domain.

*Table 7-1. Administrative commands showing optional vrouter arguments*

| <b>Command</b> | <b>Permissions</b> | <b>Usage with optional vrouter argument</b>                              |
|----------------|--------------------|--------------------------------------------------------------------------|
| arptable       | system             | arptable [vrouter] [[-a hostname MAC_address]  <br>[-d hostname]   [-f]] |

Table 7-1. Administrative commands showing optional vrouter arguments (Continued)

| Command  | Permissions | Usage with optional vrouter argument                                                                                                                                  |
|----------|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ipcache  | system      | ipcache [-r vrouter] [cache] [stats]                                                                                                                                  |
| iproute  | system      | iproute add [-r vrouter] dest_IPAddress/subnet_mask<br>gateway_IPAddress [preference] [metric]<br>iproute delete [-r vrouter] dest_IPAddress/subnet_mask<br>[gateway] |
| netstat  | system      | netstat [vrouter] [-i] [-r] [?] [-n] [-d] [-s identifiers]<br>[-z]                                                                                                    |
| vrouter  | system      | vrouter [dump [full]] [callback]                                                                                                                                      |
| nslookup | diagnostic  | nslookup [-v] [-s dnssrvr_IPAddress] [-r vrouter] hostname                                                                                                            |
| ping     | diagnostic  | ping [-q   -v] [-i delay_sec   -I delay_msec]<br>[-s packetsize] [-r vrouter] [-x source_IPAddress] hostname                                                          |
| telnet   | diagnostic  | telnet [-a   -b   -t] [-v vrouter] [-l[e]   -r[e]] hostname<br>[portnumber]                                                                                           |
| tracert  | diagnostic  | tracert [-n] [-v] [-m max_ttl] [-p port] [-q nqueries]<br>[-w waittime] [-r vrouter] [-s src_IPAddr] hostname<br>[datasize]                                           |
| info np  | debug       | info np arp [vrouter]<br>info np route [vrouter]                                                                                                                      |
| ifmgr    | debug       | ifmgr [-r vrouter] -d [ifNum] -t<br>ifmgr [up down] [ifNum ifName]                                                                                                    |

For details about the system or diagnostic commands in Table 7-1, see the *Stinger Reference*. For information about the debug-level commands, see Appendix A, “IP Control Module Diagnostics.”





---

# OSPF Configuration

# 8

|                                                               |      |
|---------------------------------------------------------------|------|
| Overview of supported OSPF features . . . . .                 | 8-1  |
| Enabling OSPF systemwide . . . . .                            | 8-9  |
| Configuring OSPF on Gigabit Ethernet . . . . .                | 8-9  |
| Configuring OSPF on an ATM trunk interface. . . . .           | 8-13 |
| Configuring global route options that apply to OSPF . . . . . | 8-16 |
| Configuring IP-ROUTE OSPF options . . . . .                   | 8-17 |
| Administrative tools for OSPF routing. . . . .                | 8-19 |

Open Shortest Path First (OSPF) is an Internet routing protocol, developed by the OSPF working group of the Internet Engineering Task Force, and defined in RFC 2328, *OSPF Version 2*.

OSPF was designed for the TCP/IP Internet environment, including explicit support for IP subnets, tagging of externally derived routing information, and use of IP multicast for sending or receiving link-state updates.

## Overview of supported OSPF features

This section provides a brief overview of OSPF routing to help you configure the system properly. For details about how OSPF works, see RFC 2328, *OSPF Version 2*. Following are the OSPF functions discussed in this section:

- Multiple authentication schemes, per RFC 1583
- Variable length subnet masks (VLSMs)
- Link state advertisement (LSA) types 1, 2, 3, 4, 5, and 7
- Backup and designated backup router capability on broadcast networks
- Nonbroadcast multiaccess (NBMA) network support over point-to-point links
- Configurable cost metrics
- Hierarchical routing via normal areas, stub areas, and not-so-stubby-areas (NSSAs)
- Shortest path first link-state routing algorithm

■ Diagnostics and traps



**Note** Stinger does not support OSPF on DSL interfaces. OSPF is supported only on ATM trunks (through trunk modules such as OC-3) and Gigabit Ethernet ports (IP2x00 only).

### **Limited border router capability**

A Stinger IP DSLAM acts as an OSPF internal router with limited border router capability. It does not currently function as an interior gateway protocol (IGP) gateway, although it performs autonomous system border router (ASBR) calculations for external routes, such as WAN links that do not support OSPF. The Stinger IP DSLAM imports external routes into its OSPF database and flags them as autonomous system external (ASE). It redistributes these routes via OSPF ASE advertisements, and propagates its OSPF routes to remote WAN routers running RIP.

### **One active IP interface per port**

If more than one IP address is assigned to the same physical port, only one of the logical interfaces can have OSPF enabled. For example, in the following listing the Gigabit Ethernet port has one virtual interface:

```
admin> dir ip-interface
27 05/06/2005 09:27:35 { { any-shelf any-slot 0 } 0 }
46 05/13/2005 10:53:46 { { shelf-1 first-control-module 1 } 0 }
39 08/12/2005 08:37:48 { { shelf-1 first-control-module 2 } 0 }
54 08/12/2005 08:37:27 { { shelf-1 first-control-module 2 } 1 }
30 08/10/2005 10:14:31 { { shelf-1 slot-14 1 } 0 }
30 08/10/2005 10:14:31 { { shelf-1 slot-14 2 } 0 }
```

OSPF can be enabled on one of the port's IP interfaces, but not on more than one interface for the same port.

### **Authentication**

All OSPF protocol exchanges are authenticated by simple authentication by default. Only trusted routers can participate in the autonomous system's routing. A variety of authentication schemes can be used. In fact, different authentication types can be configured for each area. For a discussion of areas, see "Hierarchical routing (areas)" on page 8-6.

Authentication provides added security for the routers that are on the network. Routers that do not have the password are unable to gain access to the routing information, because authentication failure prevents a router from forming adjacencies. (For a discussion of adjacencies, see "Exchange of routing information" on page 8-3.) If both sides of a connection do not support the same authentication method, packet error messages can result.

In addition to null and simple password authentication, Stinger IP DSLAMs support the MD5 cryptographic authentication method for OSPF as described RFC 2328. For details about MD5 encryption, see RFC 2328.

### **Support for variable-length subnet masks**

OSPF routers handle variable-length subnet masks (VLSMs). Each route distributed by OSPF has a destination address and subnet mask, and two different subnets of the

same IP network can use different size subnet masks. A packet is routed to the best (longest or most specific) match. Host routes are considered to be subnets whose masks are all ones (0xFFFFFFFF).



**Note** OSPF is useful for networks that use VLSMs. However, to prevent excessive link-state calculations by all OSPF routers on the network, make every effort to assign subnets that are as continuous as possible.

## Exchange of routing information

An OSPF router stores its information about the network in a topological database and propagates only changes to the database. Selected neighboring routers form relationships, referred to as *adjacencies*, for the purpose of exchanging routing information. Not every pair of neighboring routers become adjacent. Routers connected by point-to-point networks and virtual links always become adjacent. On multiaccess networks, all routers become adjacent to routers identified as the designated router (DR) and the backup designated router (BDR).

As the adjacency is established, the neighbors exchange databases and build a consistent, synchronized database between them. When an OSPF router detects a change on one of its interfaces, it modifies its topological database and multicasts the change to its adjacent neighbors, which in turn propagate the change to their adjacent neighbors, until all routers within an area have synchronized topological databases. This process provides quick convergence among routers.

A link state advertisement (LSA) is a packet that describes various aspects of an OSPF route. Each LSA is flooded throughout a routing domain. The collected LSAs of all routers and networks forms the OSPF topological database. Table 8-1 shows the types of LSAs.

Table 8-1. Description of LSA types

| LSA type                        | Description                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Type 1 (RTR) router             | Type 1 router LSAs describe the collected states of the router's interfaces.                                                                                                                                                                                                                               |
| Type 2 (NET) network            | Type 2 LSAs describe the set of routers attached to the network.                                                                                                                                                                                                                                           |
| Types 3 and 4<br>(Summary LSAs) | Summary LSAs are flooded throughout a single area. Type 3 summary LSAs describe routes to networks. Type 4 summary LSAs describe routes to autonomous system boundary routers.                                                                                                                             |
| Type 5 (ASE) AS-external        | Type 5 LSAs describe routes to destinations external to the autonomous system (AS). An AS-external-LSA can also describe a default route for the autonomous system. For example, other routers send LSAs to only the designated router by using the All-Designated-Routers multicast address of 224.0.0.6. |


Table 8-1. Description of LSA types

| LSA type          | Description                                                                                                                                                                                                                                                                                                                                                                            |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Type 7 (ASE) NSSA | NSSAs are like stub areas in that they do not receive or originate type 5 LSAs. NSSAs rely solely on default routing for external routes. They employ type 7 LSAs for carrying ASE route information within the area. Type 7 LSAs use a propagate (P) bit to flag the NSSA border router to translate the type 7 LSA into a type 5 LSA, which can then be propagated into other areas. |

## Designated and backup designated routers on broadcast networks

In OSPF terminology, a broadcast network is any network that has more than two OSPF routers attached and supports the capability to address a single physical message to all the attached routers (broadcast). Neighboring routers are discovered dynamically on these networks using the OSPF Hello protocol, which uses the broadcast capability. Ethernet is an example of a broadcast network. Figure 8-1 shows such a network.

Figure 8-1. OSPF broadcast network on Gigabit Ethernet



To reduce the number of adjacencies each router must form, OSPF calls one of the routers the designated router. As routers begin to form adjacencies, they elect a designated router and then all other routers on the network establish adjacencies, primarily with the designated router. This process simplifies the routing table update procedure and reduces the number of link-state records in the database. The designated router plays other important roles as well to reduce the overhead of OSPF link-state procedures.

To prevent the designated router from becoming a serious liability to the network if it fails, OSPF routers also elect a backup designated router at the same time. Other routers maintain adjacencies with both the designated router and its backup, but the backup router leaves as many of the processing tasks as possible to the designated router. If the designated router fails, the backup immediately becomes the designated router and a new backup is elected.

You choose the designated router on the basis of the processing power, speed, and memory of the system, then assign priorities to other routers on the network, in case both the designated and backup designated routers fail.



**Note** The Stinger IP DSLAM can function as a designated router or backup designated router. However, many sites choose to assign a LAN-based router for these roles to dedicate the Stinger IP DSLAM to WAN processing.

## Routing across NBMA interfaces

An OSPF nonbroadcast multiaccess (NBMA) network is any network that has multiple points of access (more than two routers) and does not support broadcast capability. OSPF routers operate on an NBMA network much as they do on a broadcast network, by using the Hello protocol to form adjacencies and identify the designated router. However, because the routers cannot discover their neighboring routers dynamically by means of broadcasts, you must specify some additional parameters. For Stinger IP DSLAMs, a WAN link is always point-to-point. However, you can configure the system to interact with NBMA routers via a virtual circuit across an ATM trunk interface.


## Configurable cost metrics

You assign a cost to the output side of each router interface. The lower the cost, the more likely the interface is to be used to forward data traffic. Costs can also be associated with the externally derived routing data.

The OSPF cost can also be used for preferred-path selection. If two paths to a destination have equal costs, you can assign a higher cost to one of the paths so it can be a backup to be used only when the primary path is not available.

Figure 8-2 shows how costs are used to direct traffic over high-speed links. For example, if Router-2 in Figure 8-2 receives packets destined for Host B, it routes them through Router-1 across two T1 links (cost=20) rather than across one 56Kbps B channel to Router-3 (cost=240).

Figure 8-2. OSPF costs for different types of links



The Stinger IP DSLAM has a default cost of 1 for a connected route (Ethernet) and 10 for a WAN link. If two paths have the same destination, the Stinger IP DSLAM uses the path with the lower cost unless route preferences change the equation.

When assigning costs, remember to account for the bandwidth of a connection. For example, for a single B-channel connection, the cost is 24 times greater than for a T1 link.



**Note** Be careful when assigning costs. Incorrect cost metrics can cause delays and congestion on the network.


## **Hierarchical routing (areas)**

If a network becomes too large, the size of the database, time required for route computation, and related network traffic become excessive. You can partition an autonomous system into areas to provide hierarchical routing, with a backbone area connecting the other areas. The backbone area is special and always has the area number 0.0.0.0. The backbone consists of networks not contained in any area, their attached routers, and routers that belong to multiple areas.

The backbone must be contiguous. You can use virtual links to connect two backbone routers that have an interface to a common nonbackbone area. OSPF treats two routers joined by a virtual link as if they were connected by an unnumbered point-to-point network. The backbone distributes routing information between areas and has all the properties of an area. The topology of the backbone is invisible to each of the areas, while the backbone itself has no information about area topology.

Each area acts as its own network: All area-specific routing information stays within the area, and all routers within an area must have a synchronized topological database. To tie the areas together, some routers belong to the backbone area and also to one of the other areas. These routers are area border routers (ABRs). In Figure 8-3, all the routers are ABRs.

*Figure 8-3. Dividing an OSPF autonomous system into areas*



With the ABRs and area boundaries set up correctly, link-state databases are unique to an area. You can configure the Stinger IP DSLAM to route in normal areas, stub areas, and NSSAs. These different kinds of areas handle the autonomous system external (ASE) routes originated by ASBRs in the following ways.

- **Normal areas**

An OSPF normal area allows type 5 LSAs to be flooded throughout the area.

- **Stub areas**

To reduce the cost of routing, OSPF supports stub areas in which a default route summarizes all external routes. A stub area is connected to the backbone area by one ABR, so it has only one exit point and need not maintain information about specific external routes. A stub area allows no type 5 LSAs to be propagated into or throughout the area.

- **NSSAs**

For details about the NSSA specification, see RFC 1587, *The OSPF NSSA Option*. NSSAs do not receive or originate type 5 LSAs. When a Stinger IP DSLAM is routing OSPF in an NSSA, it imports ASE routes defined in local or RADIUS

profiles as type 7 LSAs. These imported ASE LSAs have the P-bit enabled, which flags border routers to translate them into type 5 LSAs.

You can list the router IDs of NSSA border routers that are translating type 7 LSAs to type 5 LSAs, by entering the `ospf translators` command. For example:

```
admin> ospf translators
Area ID      Router ID
0.0.0.1      10.105.0.13
0.0.0.2      12.1.1.1
```

## Link-state routing algorithms

The link-state routing algorithms require that all routers within a domain maintain synchronized (identical) topological databases, and that the databases describe the complete topology of the domain. An OSPF router's domain can be an autonomous system or an area within an autonomous system.

OSPF routers create and update a link-state database from information exchanged with other routers. Link-state databases are synchronized between pairs of adjacent routers (as described in “Exchange of routing information” on page 8-3). In addition, each OSPF router uses its link-state database to calculate a self-rooted tree of shortest paths to all destinations. The routing table is built from these calculated shortest-path trees. For example, consider the network topology in Figure 8-4.

Figure 8-4. Sample OSPF topology




Table 8-2 shows the relevant information in the routers' link-state databases.

Table 8-2. Link-state databases for OSPF topology in Figure 8-4

| Router-1         | Router-2         | Router-3         |
|------------------|------------------|------------------|
| Network-1/Cost 0 | Network-2/Cost 0 | Network-3/Cost 0 |
| Network-2/Cost 0 | Network-3/Cost 0 | Network-4/Cost 0 |
| Router-2/Cost 20 | Router-1/Cost 20 | Router-2/Cost 30 |
|                  | Router-3/Cost 30 |                  |

From the link-state database, each router builds a self-rooted shortest-path tree, and then calculates a routing table stating the shortest path to each destination in the

autonomous system. (See Table 8-3, Table 8-4, and Table 8-5.) The table also includes externally derived routing information.

All the routers calculate a routing table of shortest paths, based on the link-state database. Externally derived routing data is advertised throughout the autonomous system but is kept separate from the link-state data. Each external route can also be tagged by the advertising router, enabling the passing of additional information between routers on the boundary of the autonomous system.

Table 8-3. Shortest-path tree and resulting routing table for Router-1


|                                                                                   |                    |                 |               |
|-----------------------------------------------------------------------------------|--------------------|-----------------|---------------|
|  | <i>Destination</i> | <i>Next hop</i> | <i>Metric</i> |
|                                                                                   | Network-1          | Direct          | 0             |
|                                                                                   | Network-2          | Direct          | 0             |
|                                                                                   | Network-3          | Router-2        | 20            |
|                                                                                   | Network-4          | Router-2        | 50            |

Table 8-4. Shortest-path tree and resulting routing table for Router-2



|                                                                                     |                    |                 |               |
|-------------------------------------------------------------------------------------|--------------------|-----------------|---------------|
|  | <i>Destination</i> | <i>Next hop</i> | <i>Metric</i> |
|                                                                                     | Network-1          | Router-1        | 20            |
|                                                                                     | Network-2          | Direct          | 0             |
|                                                                                     | Network-3          | Direct          | 0             |
|                                                                                     | Network-4          | Router-3        | 30            |

Table 8-5. Shortest-path tree and resulting routing table for Router-3

|                                                                                     |                    |                 |               |
|-------------------------------------------------------------------------------------|--------------------|-----------------|---------------|
|  | <i>Destination</i> | <i>Next hop</i> | <i>Metric</i> |
|                                                                                     | Network-1          | Router-2        | 50            |
|                                                                                     | Network-2          | Router-2        | 30            |
|                                                                                     | Network-3          | Direct          | 0             |
|                                                                                     | Network-4          | Direct          | 0             |



## Enabling OSPF systemwide

Before the Stinger system can route OSPF, it must be configured for IP routing. For details about configuring IP routing, see Chapter 4, “IP Routing Configuration.”

To configure the system to use OSPF routing, you must configure each LAN or WAN interface that will support OSPF routing, and enable the protocol systemwide.

The following parameters, shown with default values, enable the protocol and specify global settings:

```
[in IP-GLOBAL:ospf-global]
enable = no
as-boundary-router = yes
ospf-max-lsa = 0
```

| Parameter          | Setting                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| enable             | Enables or disables the OSPF protocol systemwide. Set to <b>yes</b> to enable the protocol. If set to <b>no</b> (the default), the protocol is disabled systemwide. If you are modifying several OSPF-related profiles, you can use the <b>no</b> setting to prevent the OSPF subsystem from reinitializing whenever you write a modified profile. Then set the parameter to <b>yes</b> when the modifications are complete. A change to the setting takes effect immediately after you write the profile. |
| as-boundary-router | Enables or disables ASBR calculations related to external routes. If this parameter is set to <b>yes</b> , the system originates type 5 AS-external LSAs for the IP address associated with IP-ROUTE profiles, and for IP-INTERFACE and CONNECTION profiles for which OSPF is <i>not</i> enabled. (Of course, for interfaces with OSPF enabled, the system originates type 1 router LSAs.)                                                                                                                 |
| ospf-max-lsa       | Maximum number of LSAs allowed in the link-state database. Specify a number from 0 through 4,294,967,295. The default setting is 0.                                                                                                                                                                                                                                                                                                                                                                        |

For example, the following commands enable the OSPF protocol:

```
admin> read ip-global
admin> set ospf-global enable = yes
admin> write -f
```

## Configuring OSPF on Gigabit Ethernet

Before the Stinger system can route OSPF, it must be configured for IP routing. For details about configuring IP-INTERFACE profiles, see “Configuring IP-INTERFACE profiles for Ethernet ports” on page 4-6.

## Overview of IP-INTERFACE OSPF settings

Following are the OSPF parameters, shown with default values, for configuring OSPF routing on a GigE interface:

```
[in IP-INTERFACE/{ { any-shelf any-slot 0 } 0 }:ospf]
active = no
area = 0.0.0.0
area-type = normal
hello-interval = 10
dead-interval = 40
priority = 5
authen-type = simple
auth-key = *****
key-id = 0
cost = 1
down-cost = 16777215
ase-type = type-1
ase-tag = c0:00:00:00
transit-delay = 1
retransmit-interval = 5
non-multicast = no
network-type = Broadcast
poll-interval = 10
profile-type = lan
md5-auth-key = *****
```

| Parameter      | Setting                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| active         | Enables or disables OSPF on an interface.                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| area           | OSPF area number in dotted decimal notation. The default 0.0.0.0 represents the backbone area.                                                                                                                                                                                                                                                                                                                                                                                |
| area-type      | Type of area. The default is the <b>normal</b> area type.                                                                                                                                                                                                                                                                                                                                                                                                                     |
| hello-interval | Number of seconds between Hello packets.                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| dead-interval  | Number of seconds without receiving a Hello packet before instituting a link-state change.                                                                                                                                                                                                                                                                                                                                                                                    |
| priority       | Priority value, from 0 to 255, used to elect a designated router and backup designated router. A setting of 0 excludes the Stinger IP DSLAM from becoming a designated router or backup designated router. The higher the priority value of the Stinger system relative to other OSPF routers on the network, the better the chances that it will become one of these routers. For details, see “Designated and backup designated routers on broadcast networks” on page 8-4. |
| authen-type    | Type of authentication to use.<br><div>none                      No authentication is required.</div>                                                                                                                                                                                                                                                                                                                                                                         |


| Parameter                  | Setting                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                            | <p><b>simple</b>      The router uses the password supplied in the <b>auth-key</b> parameter to validate OSPF packet exchanges. This is the default value.</p> <p><b>md5</b>      The router uses MD5 encryption and the authentication key ID supplied by the <b>key-id</b> parameter to validate OSPF packet exchanges. For related information, see “Authentication” on page 8-2.</p> |
| <b>auth-key</b>            | Secret key for authenticating traffic in the router’s area. Enter a text string of up to 8 characters. When <b>authen-type</b> is set to <b>md5</b> , you must set the <b>md5-auth-key</b> parameter to specify a key.                                                                                                                                                                   |
| <b>key-id</b>              | Number from 0 to 255, used to encrypt the secret key when <b>authen-type</b> is set to <b>md5</b> .                                                                                                                                                                                                                                                                                      |
| <b>cost</b>                | Cost of routing to the interface. The lower the cost assigned to a route, the more likely it is to be used to forward traffic. For details, see “Configurable cost metrics” on page 8-5.                                                                                                                                                                                                 |
| <b>down-cost</b>           | Cost applied to the interface when it is unavailable. The output cost when the link is physically unavailable but virtually active.                                                                                                                                                                                                                                                      |
| <b>ase-type</b>            | Type of metric to apply to routes learned from RIP. This parameter applies in a CONNECTION profile only when OSPF is <i>not</i> active. <p><b>type-1</b>      Expresses the metric in the same units as the interface cost. This is the default.</p> <p><b>type-2</b>      Metric is larger than any link-state path.</p>                                                                |
| <b>ase-tag</b>             | If OSPF is disabled on this interface (if the <b>active</b> parameter is set to <b>no</b> in the subprofile), and the system is configured to perform ASBR calculations (as described in “Enabling OSPF systemwide” on page 8-9), the system originates AS-external LSAs for the interface IP address and uses this tag to populate the External Route Tag field of the LSA.             |
| <b>transit-delay</b>       | Estimated number of seconds required to transmit a Link State Update packet over this interface, taking into account transmission and propagation delays. On a connected route, you can leave the default of 1.                                                                                                                                                                          |
| <b>retransmit-interval</b> | Number of seconds between LSA retransmissions for adjacencies belonging to this interface. Its value is also used when retransmitting database description and link-state request packets. On a typical connected route, accept the default setting of 5.                                                                                                                                |

| Parameter     | Setting                                                                                                                                                                                                                               |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| non-multicast | Supports NBMA configuration to a GRF® multigigabit router. See “Sample configuration of NBMA across point-to-point” on page 8-14.                                                                                                     |
| network-type  | Type of OSPF interface.<br>broadcast      A broadcast-capable network, such as Ethernet.<br>nonbroadcast    An NBMA network, such as a trunk interface.<br>point-to-point   A point-to-point network, consisting of two routers only. |
| poll-interval | Not used on a broadcast network.                                                                                                                                                                                                      |
| profile-type  | A read-only parameter used internally to verify settings in the profile.                                                                                                                                                              |
| md5-auth-key  | Secret key to be used for MD5 authentication, up to 16 characters. The default value is <code>ascend0</code> . When <code>authen-type</code> is set to <code>md5</code> , you must supply a value for this parameter.                 |

## Sample Gigabit Ethernet interface configuration

Figure 8-5 shows three OSPF routers in the backbone area of an autonomous system. Because all OSPF routers are in the same area, the routers form adjacencies and synchronize their databases. This example shows how to configure the Gigabit Ethernet interface of the unit labeled Stinger-2.

Figure 8-5. OSPF on a LAN interface



All OSPF routers in Figure 8-5 have RIP turned off. Running both RIP and OSPF is unnecessary, and turning RIP off reduces processor overhead. OSPF can learn routes from RIP interfaces, incorporate them in the routing table, assign them an external metric, and tag them as external routes.

Although RFC 2328 does not specify a limitation for the number of routers in the backbone area, keep the number of routers relatively small, because changes that occur in area zero are propagated throughout the autonomous system. Another way to configure the same units is to create a second area (such as 0.0.0.1) in one of the

existing OSPF routers, and add the Stinger IP DSLAM to that area. You can then assign the same area number (0.0.0.1) to all OSPF routers reached through the Stinger IP DSLAM across a WAN link.

The following sample commands show how to configure Stinger-2 in Figure 8-5. The commands assign the IP address 10.168.8.17/24 to the local interface and configure the OSPF router in the backbone area:

```
admin> read ip-interface { { 1 8 2 } 0 }
admin> set ip-address = 10.168.8.17/24
admin> set rip-mode = routing-off
admin> set ospf active = yes
admin> write -f
```

The following sample commands show how to configure the IP interface for MD5 authentication:

```
admin> read ip-interface { { 1 8 2 } 0 }
admin> set ospf authen-type = md5
admin> set ospf md5-auth-key = 12!secret*34key
admin> write -f
```

## **Configuring OSPF on an ATM trunk interface**

Before the Stinger IP DSLAM can route OSPF, it must be configured for IP routing. For details about configuring CONNECTION profiles for IP routing, see “Configuring IPoA subscriber connections” on page 4-49. In Stinger systems, OSPF routing across an ATM trunk interface uses a point-to-point link.

### **Overview of CONNECTION OSPF settings**

Following are the parameters, shown with default settings, for configuring OSPF on an ATM trunk interface.

```
[in CONNECTION/"" :ip-options:ospf-options]
active = no
area = 0.0.0.0
area-type = normal
hello-interval = 30
dead-interval = 120
priority = 5
authen-type = simple
auth-key = *****
key-id = 0
cost = 10
down-cost = 1000
ase-type = type-1
ase-tag = c0:00:00:00
transit-delay = 1
retransmit-interval = 5
non-multicast = no
network-type = Point-to-Point
poll-interval = 10
```

## OSPF Configuration

Configuring OSPF on an ATM trunk interface

---


```
profile-type = wan
md5-auth-key = *****
```

These are the same parameters described for enabling OSPF on the Gigabit Ethernet interface. For definitions, see “Overview of IP-INTERFACE OSPF settings” on page 8-10.

### Sample OSPF point-to-point configuration

This example shows how to configure a CONNECTION profile in the system labeled Stinger-2 in Figure 8-6, to enable it to route OSPF across the ATM cloud to Stinger-1. In this example, the unit labeled Stinger-1 uses the IP address 3.3.3.3/24.

Figure 8-6. OSPF over ATM point to point



The following commands configure OSPF in the unit labeled Stinger-2 in Figure 8-6:

```
admin> read conn stinger1-atmvc
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 3.3.3.3/24
admin> set ip-options rip = routing-off
admin> set ip-options ospf active = yes
admin> set atm-options atm1483type = aal5-llc
admin> set atm-options vci = 101
admin> set atm-options nailed-group = 802
admin> write -f
```

### Sample configuration of NBMA across point-to-point

With the current software version, NBMA is supported only on a point-to-point WAN link. Full support for multiaccess is not provided.

#### Overview of additional NBMA settings

In addition to the standard settings for OSPF point-to-point, the following parameters, shown with default settings, must be configured to support NBMA:


```
[in CONNECTION/"" :ip-options:ospf-options]
non-multicast = no
network-type = Point-to-Point
[in CONNECTION/"" :ip-options]
local-address = 0.0.0.0/0
[in OSPF-NBMA-NEIGHBOR/"" ]
name* = ""
host-name = ""
ip-address = 0.0.0.0
dr-capable = no
```

| Parameter                         | Setting                                                                                                                                                                                                                                                                     |
|-----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| non-multicast                     | For an NBMA connection to a GRF® multigigabit router, the <b>non-multicast</b> parameter must be set to <b>yes</b> . This causes the translation of the multicast traffic to directed traffic. This setting is required only when connecting to a GRF® multigigabit router. |
| network-type                      | For NBMA, the <b>network-type</b> parameter must be set to <b>NonBroadcast</b> .                                                                                                                                                                                            |
| local-ip-address                  | When <b>network-type</b> is set to <b>NonBroadcast</b> , a <b>local-ip-address</b> value must be provided or the system reports a configuration error such as the following:<br><br>OSPF CONFIG ERROR: NBMA profile router-2 has illegal 0.0.0.0 address                    |
| ospf-nbma-neighbor:<br>name       | Name of the <b>ospf-nbma-neighbor</b> profile.                                                                                                                                                                                                                              |
| ospf-nbma-neighbor:<br>host-name  | Name of the local <b>CONNECTION</b> profile that defines the connection to the neighboring router.                                                                                                                                                                          |
| ospf-nbma-neighbor:<br>ip-address | IP address of the neighboring router.                                                                                                                                                                                                                                       |
| ospf-nbma-neighbor:<br>dr-capable | Whether the neighboring router can be the designated router (yes or no).                                                                                                                                                                                                    |

## Example of an NBMA configuration

Figure 8-7 shows a Stinger IP DSLAM connecting point-to-point with another router that operates in an OSPF NBMA network.

*Figure 8-7. OSPF NBMA over ATM point to point*



## OSPF Configuration

### Configuring global route options that apply to OSPF

---

Following is a CONNECTION profile on the system labeled Stinger-1 to enable NBMA on the link to Router-2:

```
admin> new connection router-2
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 90.90.90.2/8
admin> set ip-options local-address = 90.90.90.1/8
admin> set ip-options ospf-options active = yes
admin> set ip-options ospf-options authen-type = simple
admin> set ip-options ospf-options auth-key = mykey
admin> set ip-options ospf-options network-type = NonBroadcast
admin> set atm-options atm1483type = aal5-llc
admin> set atm-options vci = 101
admin> set atm-options nailed-group = 851
admin> write -f
```

The following profile enables the unit to form an adjacency with Router-2:

```
admin> new ospf-nbma-neighbor router-2
admin> set host-name = router-2
admin> set ip-address = 90.90.90.2/8
admin> write -f
```

## Configuring global route options that apply to OSPF

The IP-GLOBAL profile contains several settings that apply only when OSPF routing is in use. Following are the relevant parameters, shown here with their default settings:

```
[in IP-GLOBAL]
pool-ospf-adv-type = type-1
ospf-pref = 10
ospf-ase-pref = 150
ospf-global = { no yes 0 }
rip-tag = c8:00:00:00
rip-ase-type = 1
```

| Parameter          | Setting                                                                                                                                                                                                                            |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| pool-ospf-adv-type | Type of ASE metric applied to summarized pools imported into OSPF as external routes.                                                                                                                                              |
| ospf-pref          | Preference value for routes learned from OSPF. Valid values are 0 to 255. The default value is 10.                                                                                                                                 |
| ospf-ase-pref      | Preference value for routes learned from RIP, ICMP, or another non-OSPF protocol. Specify a value from 0 through 255. By default, routes learned dynamically from another routing protocol are assigned a preference value of 150. |



| Parameter    | Setting                                                                                                             |
|--------------|---------------------------------------------------------------------------------------------------------------------|
| rip-tag      | Hexadecimal number associated with routes learned from RIP. OSPF border routers can use the tag to filter a record. |
| rip-ase-type | Type of ASE metric applied to routes learned from RIP.                                                              |

### **Example of importing a summarized pool as an ASE**

The following commands configure a summarized pool and import it to OSPF with a type 1 OSPF metric:

```
admin> read ip-global
admin> set pool-summary = yes
admin> set pool-base-address 1 = 10.12.253.1
admin> set assign-count 1 = 62
admin> set pool-ospf-adv-type = type-1
admin> write -f
```

When pool-summary is set to yes and OSPF is enabled, the OSPF subsystem uses the pool-ospf-adv-type parameter to determine how to import summarized routes into OSPF. If this parameter is set to **type-1**, the metric for the route to a summarized pool is expressed in the same units as the link-state metric (interface cost).

If pool-ospf-adv-type is set to **type-2**, the unit considers the routing between autonomous systems as the major cost of routing a packet, and conversion of external costs to internal link-state metrics is unnecessary. If the parameter is set to **internal**, the summarized pool addresses are imported into OSPF as intra-area routes, which enables them to work properly with stub areas.

### **Example of setting ASE preferences**

The ospf-pref and ospf-ase-pref parameters determine the preference values assigned to routes learned from other OSPF routers and those imported from other dynamic routing protocols. The default settings place a much lower preference on OSPF routes, which means that the routes learned from other protocols (ASE routes) are more likely to be used. The following commands decrease to 100 the preference assigned to ASE routes (the default is 150):

```
admin> read ip-global
admin> set ospf-ase-pref = 100
admin> write -f
```

## **Configuring IP-ROUTE OSPF options**

For details about configuring static routes, see “Configuring IP-ROUTE profiles” on page 4-40. The following parameters in the IP-ROUTE profile (shown with sample settings), apply only when OSPF is enabled:

```
in IP-ROUTE/[""]
cost = 1
ase-type = type-1
ase-tag = c0:00:00:00
```

ase7-adv = N/A

| Parameter | Setting                                                                                                                                                                                                                                        |
|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| cost      | Cost of routing to the interface. The lower the cost, the more likely the interface is to be used to forward traffic. See “Configurable cost metrics” on page 8-5.                                                                             |
| ase-type  | Type of metric to apply to routes learned from RIP. The default value of <b>type-1</b> expresses the metric in the same units as the interface cost. With the value of <b>type-2</b> , the metric is larger than any link-state path.          |
| ase-tag   | If the system is configured to perform ASBR calculations (as described in “Enabling OSPF systemwide” on page 8-9), the system originates AS-external LSAs for the route and uses this tag to populate the External Route Tag field of the LSA. |
| ase7-adv  | <i>Currently not used.</i>                                                                                                                                                                                                                     |

### **Example of configuring a type 7 LSA in an NSSA**

For background information about NSSAs, see “Hierarchical routing (areas)” on page 8-6. To configure the Stinger IP DSLAM to route OSPF in an NSSA, *all* OSPF interfaces in the Stinger system must specify the NSSA area-type.

To configure a type 7 LSA, you must specify a static route in an IP-ROUTE profile. Following are the related parameters (shown with sample settings):

```
[in IP-ROUTE/external]
name* = external
dest-address = 10.4.5.0/22
gateway-address = 10.4.5.7
metric = 0
cost = 1
preference = 100
third-party = no
ase-type = type-1
ase-tag = c0:00:00:00
private-route = no
active-route = yes
ase7-adv = n/a
```

The following procedure configures the Stinger IP DSLAM to route in an NSSA and imports a type 7 LSA that specifies an external route across the WAN link:

- 1 Assign an NSSA area type to the IP interface that is running OSPF. For example:  
admin> read ip-interface { { 1 8 2 } 0 }  
admin> set ospf area-type = nssa  
admin> write -f
- 2 Configure the WAN link that represents an ASE route. For example:  
admin> read connection ase-link  
admin> set ip-options remote = 10.4.5.7/22

```
admin> set ip-options rip = routing-off
admin> set ip-options ospf active = yes
admin> write -f
```

- 3 Configure a static route to the remote site. For example:

```
admin> new ip-route type7
admin> set dest = 10.4.5.0/22
admin> set gateway = 10.4.5.7
admin> write -f
```

### ***Example of assigning a cost to a static route***

The lower the cost assigned to a route, the more likely the router is to choose the route to forward traffic. Typically, you account for the bandwidth of a connection when assigning costs. The Stinger IP DSLAM has a default cost of **1** for a connected route (Ethernet) and **10** for an ATM VC. If two paths have the same destination, the Stinger IP DSLAM uses the path with the lower cost.



**Note** Be careful when assigning costs. Incorrect cost metrics can cause delays and congestion on the network.

In the following example, an administrator assigns a cost of 25 to a static route:

```
admin> new ip-route myLink
admin> set dest = 10.1.2.0/24
admin> set gateway = 10.9.8.10
admin> set cost = 25
admin> write -f
```

### ***Administrative tools for OSPF routing***

The `ospf` diagnostic-level commands display information related to OSPF routing, including the link state advertisements (LSAs); the routing table for border routers; and the OSPF areas, interfaces, statistics, and routing table. For details, see the chapter on monitoring OSPF in the *Stinger Administration Guide*.



# IP Multicast Configuration

# 9


|                                                        |      |
|--------------------------------------------------------|------|
| IP multicast forwarding .....                          | 9-1  |
| Configuring MBONE interfaces .....                     | 9-3  |
| Managing multicast group memberships .....             | 9-8  |
| Configuring multicast client interfaces .....          | 9-11 |
| Multicast server virtual circuits .....                | 9-22 |
| Distributed multicast with HB LIMs (IP2100 only) ..... | 9-25 |
| Administrative tools for IGMP operations .....         | 9-31 |
| CAC for multicast channels .....                       | 9-31 |
| Configuring multicast CAC on Stinger CM .....          | 9-38 |
| Viewing multicast CAC statistics .....                 | 9-39 |

## IP multicast forwarding

IP multicast forwarding enables a Stinger IP DSLAM to receive multicast transmissions from multicast backbone (MBONE) routers and forward the transmissions to multiple client interfaces.

A common use for IP multicast is to transmit streaming video across the Internet to applications running on multiple PCs or set-top boxes (STB) for television sets, as shown in Figure 9-1.

Figure 9-1. Multicast video sample setup



To the multicast clients, the Stinger IP DSLAM appears to be a multicast router originating the video stream. To the multicast routers, the Stinger IP DSLAM appears to be a multicast client, initiating and responding to group management messages via Internet Group Management Protocol (IGMP) version-1 or version-2.

To receive a transmission, the client interfaces must join a specific multicast group. A *multicast group* is a Class D IP address (from 224.0.0.0 to 239.255.255.255). When data is sent to an address in that range, it is multicast to all hosts that have joined that group. The Stinger forwards IGMP messages between clients (hosts) and the multicast router to enable these transactions.

## **Network-side MBONE interfaces**


An MBONE interface, used to receive transmissions from a multicast router, can be either a WAN MBONE on an ATM trunk interface or a LAN MBONE on a Gigabit Ethernet interface. On an MBONE interface, the Stinger IP DSLAM responds to multicast router queries, and returns responses from clients to the router.

The global IP router can receive multicast data from up to 32 MBONE interfaces. Multiple MBONE capability is not currently supported for virtual routers.

The multiple MBONE interfaces can be WAN or LAN, or any combination of the two interface types, but the total number of MBONE interfaces cannot exceed 32.

Figure 9-2 shows a Stinger IP DSLAM with three MBONE interfaces:

*Figure 9-2. Multiple MBONE interfaces on trunk or LAN interfaces*



## **Notice about Gigabit Ethernet redundancy for a LAN MBONE**

For Stinger systems with redundant IP control modules, you can configure Gigabit Ethernet redundancy for a LAN MBONE interface to enable the system to maintain MBONE operations across a controller switchover. For details, see “Configuring a redundant LAN MBONE” on page 2-8.

## **LIM-side multicast client interfaces**

The multicast clients must be on LIM interfaces, accessing the Stinger IP DSLAM through an ATM virtual circuit.

Transmission of multicast data to multicast clients on a local Ethernet interface is not currently supported. For that reason, the following parameters not currently used in the IP-INTERFACE profile and its *igmp-option* subprofile:

---

#### Unused IP-INTERFACE settings

---

```
[in IP-INTERFACE/{ { any-shelf any-slot 0 } 0 }]
multicast-rate-limit = 100
multicast-group-leave-delay = 0
multicast-group-leave-delay-msec = 0
multicast-service-profile = ""
multicast-max-groups = 0

[in IP-INTERFACE/{ { any-shelf any-slot 0 } 0 }:igmp-options]
robust-count = 2
query-interval = 125
query-response-interval = 100
last-member-query-interval = 10
last-member-query-count = 2
```

---

### ***IP2100 and IP2000 maximum limits on multicast and bridge groups***

Table 9-1 shows limits that apply with the IP2100 and IP2000 control modules in standalone systems:

*Table 9-1. Maximum multicast groups, multicast users, and bridge groups*

| IP2100 control module      |                        | IP2000 control module     |       |
|----------------------------|------------------------|---------------------------|-------|
| Multicast groups           | 512                    | Multicast groups          | 512   |
| Multicast users per group: |                        | Multicast users per group | 1,017 |
| with centralized multicast | 1,017                  |                           |       |
| with distributed multicast | 3,024<br>(theoretical) |                           |       |
| Bridge groups              | 42                     | Bridge groups             | 63    |

Because system resources are shared between bridge groups, multicast proxy, and IGMP snooping connections, the upper limit on multicast groups and bridge groups varies depending on the overall configuration of these features. For example, if bridge groups are configured in the system, the practical limit on multicast groups is decreased by the number of configured bridge groups. The number of users within a group also affects the maximum limits for multicast and bridge group configurations.

## ***Configuring MBONE interfaces***

To configure an MBONE interface, you must complete the following steps:

- 1 Enable multicast forwarding in the IP-GLOBAL profile.
- 2 Specify a profile index for each MBONE interface in the IP-GLOBAL profile.
- 3 Configure the IP-INTERFACE or CONNECTION profile for each LAN or WAN MBONE interface. Be sure to set `multicast-allowed` to `yes`.

A Stinger IP DSLAM does not support multicast heartbeat monitoring, so the following IP-GLOBAL settings are not used:

---

#### Unused IP-GLOBAL settings

---

```
[in IP-GLOBAL]
multicast-hbeat-addr = 0.0.0.0
multicast-hbeat-port = 0
multicast-hbeat-slot-time = 0
multicast-hbeat-Number-Slot = 0
multicast-hbeat-Alarm-threshold = 0
multicast-hbeat-src-addr = 0.0.0.0
multicast-hbeat-src-addr-mask = 0.0.0.0
```

---

If you need more information about these settings, see the parameter descriptions in the *Stinger Reference*.

## Overview of multiple MBONE configuration

The following parameters, shown with their default settings, are used to specify from MBONE interfaces:

```
[in IP-GLOBAL]
multicast-forwarding = no

[in IP-GLOBAL:multiple-mbone:mbone-profile]
mbone-profile[1] = ""
mbone-profile[2] = ""
mbone-profile[3] = ""
mbone-profile[4] = ""
...
mbone-profile[32] = ""

[in IP-GLOBAL:multiple-mbone:mbone-lan-interface]
mbone-lan-interface[1] = { { any-shelf any-slot 0 } 0 }
mbone-lan-interface[2] = { { any-shelf any-slot 0 } 0 }
mbone-lan-interface[3] = { { any-shelf any-slot 0 } 0 }
mbone-lan-interface[4] = { { any-shelf any-slot 0 } 0 }
...
mbone-lan-interface[32] = { { any-shelf any-slot 0 } 0 }

[in IP-INTERFACE/{ { shelf-1 first-control-module 2 } 0 }]
multicast-allowed = no

[in CONNECTION/"" :ip-options]
multicast-allowed = no
```

| Parameter            | Setting                                                                                                                                                                                                           |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| multicast-forwarding | Enables or disables multicast forwarding. When you change the value to <b>yes</b> and write the profile, the multicast subsystem reads the values in the IP-GLOBAL profile and initiates the forwarding function. |




| Parameter                           | Setting                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>mbone-profile[N]</code>       | Array of 32 indexed parameters for specifying the name of a local CONNECTION profile that provides access to an MBONE router across a trunk interface. This configures up to 32 WAN MBONE interfaces across trunk ports. The total number of MBONE interfaces specified in either these parameters or the <code>mbone-lan-interface</code> parameters, or both, cannot exceed 32.                                                                                                                                                                    |
| <code>mbone-lan-interface[N]</code> | Array of 32 indexed parameters for specifying the index of an IP-INTERFACE profile that provides access to an MBONE router across an Ethernet interface. This configures up to 32 LAN MBONE interfaces. The total number of MBONE interfaces specified in either these parameters or the <code>mbone-profile</code> parameters, or both, cannot exceed 32.<br><br>For systems with redundant controllers, you can configure Gigabit Ethernet redundancy for a LAN MBONE interface. For details, see “Configuring a redundant LAN MBONE” on page 2-8. |
| <code>multicast-allowed</code>      | Enables or disables handling of IGMP requests and responses on the LAN (IP-INTERFACE) or trunk (CONNECTION) MBONE interface.                                                                                                                                                                                                                                                                                                                                                                                                                         |

### Sample configuration with multiple MBONE interfaces

The sample setup in Figure 9-3 shows three MBONE interfaces, one on the Gigabit Ethernet interface (`{ { 1 8 2 } 0 }`) and two additional MBONE interfaces on the two trunk interfaces in slot 17.

Figure 9-3. Sample configuration with multiple MBONE interfaces



The following commands configure a LAN MBONE interface on a Gigabit Ethernet port of the IP control module:

```
admin> read ip-interface { { 1 8 2 } 0 }
admin> set ip-address = 1.1.1.2/28
admin> set multicast-allowed = yes
admin> write -f
```

The next commands configure two WAN MBONE interfaces on ATM trunk interfaces:

```
admin> new connection mcast1-17-1
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 3.3.3.3/29
admin> set ip-options multicast-allowed = yes
admin> set atm-options vci = 101
admin> set atm-options nailed-group = 801
admin write -f
admin> new connection mcast1-17-2
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 4.4.4.4/29
admin> set ip-options multicast-allowed = yes
admin> set atm-options vci = 101
admin> set atm-options nailed-group = 802
admin write -f
```

The following commands enable the multicast forwarding function and specify the MBONE interfaces:


```
admin> read ip-global
admin> set multicast-forwarding = yes
admin> set multiple-mbone mbone-profile 1 = mcast1-17-1
admin> set multiple-mbone mbone-profile 2 = mcast1-17-2
admin> set multiple-mbone mbone-lan-interface 1 = { { 1 8 2 } 0 }
admin> write -f
```

### **Sample MBONE configuration on Gigabit Ethernet VLANs**

You can set up a virtual IP interface on the Gigabit Ethernet port as a VLAN to be used for routing purposes as well as for handling multicast traffic. For information about the VLAN configuration, see “Configuring routed VLANs” on page 3-35.

Figure 9-4 shows a basic representation of this type of configuration, which enables the system to handle a multicast data stream with a VLAN tag, and forward the multicast data stream to multicast clients.

Figure 9-4. Sample configuration of VLAN MBONE interface



Following is a sample Gigabit Ethernet VLAN MBONE interface configuration.

- 1 Enable bridging on the physical interface.  

```
admin> read ethernet { 1 8 2 }
```

```
admin> set bridging-enabled = yes
```

```
admin> write -f
```
  - 2 Enable multicast forwarding in the IP-GLOBAL profile.  

```
admin> read ip-global
```

```
admin> set multicast-forwarding = yes
```

```
admin> write -f
```
  - 3 Create a VLAN-ETHERNET profile. This example uses VLAN ID 101.  

```
admin> new vlan-ethernet { { 1 8 2 } 101 }
```

```
admin> set enabled = yes
```

```
admin> write -f
```
  - 4 Create an IP-INTERFACE profile with both multicast and VLAN enabled. Specify the same VLAN ID used in the VLAN-ETHERNET profile (101 in this example).  

```
admin> new ip-interface { { 1 8 2 } 1 }
```

```
admin> set ip-address = 13.13.13.13/24
```

```
admin> set multicast-allowed = yes
```

```
admin> set vlan-enabled = yes
```

```
admin> set vlan-id = 101
```

```
admin> write -f
```
-  **Note** If you enable multicast without enabling VLAN, the system refuses to write the profile and displays the following error message:
- Multicast can be enabled on a Logical IP interface only when VLAN is enabled
- 5 Configure the new multicast and VLAN enabled interface as an MBONE interface in the IP-GLOBAL profile.  

```
admin> read ip-global
```

```
admin> set multiple-mbone mbone-lan 1 physical-address shelf = 1
```

```
admin> set multiple-mbone mbone-lan 1 physical-address slot = 8
```

```
admin> set multiple-mbone mbone-lan 1 physical-address item = 2
admin> set multiple-mbone mbone-lan 1 logical-item = 1
admin> write -f
```

## Managing multicast group memberships

To receive a multicast transmission, a client interface must join a specific multicast group. A multicast group is a class D IP address range from 224.0.0.0 to 239.255.255.255. Multicast groups from 224.0.0.100 through 224.0.0.160 are reserved for internal use. When data is sent to a valid multicast group, it is multicast to all hosts that have joined that group.

The MCAST-SERVICE profile provides a way to manage which multicast groups can be accessed by a client interface. Each profile specifies a number of multicast groups (up to 256), and a filter type. The filter type determines how the list of multicast groups is used: to allow access only to those groups, or allow access to all groups *except* those listed.

You can configure multiple MCAST-SERVICE profiles, one for each level of multicast services you provide. You can define the profiles locally or via RADIUS. The number of active MCAST-SERVICE profiles that can be provisioned is limited only by memory considerations in the system.

For examples of applying MCAST-SERVICE profiles to a client interface, see “Configuring multicast client interfaces” on page 9-11.

## Overview of MCAST-SERVICE settings

Following are the MCAST-SERVICE settings, shown with default values:

```
[in MCAST-SERVICE/""]  
service-name* = ""  
active = no  
filter-type = none  
filter-list = [ { no 0.0.0.0 1 } { no 0.0.0.0 1 } { no 0.0.0.0 1 } { no 0.0.0.+  
[in MCAST-SERVICE/"" :filter-list[1]]  
active = no  
mcast-ip-address = 0.0.0.0  
group-range-count = 1
```

| Parameter    | Setting                                                                                                                                                                                                                                                                                                                                           |
|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| service-name | Name assigned to the service profile, up to 31 characters.                                                                                                                                                                                                                                                                                        |
| active       | Enables or disables the profile. If the profile is disabled, none of the clients on the interfaces whose profiles point to this service profile are allowed to join any multicast groups. If you want the opposite effect, to allow clients on those interfaces access to all multicast groups, use the <code>filter-type</code> setting instead. |

| Parameter                            | Setting                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| filter-type                          | <p>Specifies whether access to the multicast groups defined in the filter list will be filtered inclusively or exclusively.</p> <p>With inclusive filtering, client interfaces have access only to those groups specified in the filter list.</p> <p>Currently, exclusive filters (to allow access to all groups <i>except</i> those specified in the filter list) is not supported.</p> <p>If you set this parameter to <code>none</code>, access to all multicast groups is allowed.</p>                                                                                                                                                                                                |
| filter-list[n]:<br>active            | <p>The filter list contains 256 indexed subprofiles, each of which specifies a multicast group address filter. The <code>active</code> parameter enables or disables the filter. When the filter is enabled, access to the address or address range specified in the filter is controlled on the basis of the <code>filter-type</code> setting. If the filter is disabled, it is not used.</p>                                                                                                                                                                                                                                                                                            |
| filter-list[n]:<br>mcast-ip-address  | <p>Class D IP address from 224.0.0.0 to 239.255.255.254. When data is sent to this address, it is multicast to all hosts that have joined that group. The address 239.255.255.255 cannot currently be used as a valid multicast address in a filter rule.</p>                                                                                                                                                                                                                                                                                                                                                                                                                             |
| filter-list[n]:<br>group-range-count | <p>The number of group addresses in a range of contiguous addresses that begins with the value of the <code>mcast-ip-address</code> parameter. Valid values are from 1 to 1024. The default value is 1, which indicates the single address specified as the <code>mcast-ip-address</code>.</p> <p>For example, if <code>mcast-ip-address</code> is set to 234.1.1.1 and <code>group-range-count</code> is set to 5, the filter affects addresses 234.1.1.1 through 234.1.1.5. Wraparound is permitted; for example, if <code>mcast-ip-address</code> is set to 234.1.1.230 and <code>group-range-count</code> is set to 37, the filter affects groups 234.1.1.230 through 234.1.2.10.</p> |

### Sample multicast address filters

In this example, a multicast video server supports two multicast group addresses. Multicast clients can subscribe to either the “bronze” or the “gold” multicast service. Bronze service permits access to the group at 239.255.129.119. Gold service permits access to both 239.255.129.119 and a premium group at 239.255.129.120.

The following commands configure the MCAST-SERVICE profiles:

```
admin> new mcast-service bronze-service
admin> set active = yes
admin> set filter-type = inclusive
admin> set filter-list 1 active = yes
admin> set filter-list 1 mcast-ip-address = 239.255.129.119
admin> write -f
```

```
admin> new mcast-service gold-service
admin> set active = yes
admin> set filter-type = inclusive
admin> set filter-list 1 active = yes
admin> set filter-list 1 mcast-ip-address = 239.225.129.119
admin> set filter-list 2 active = yes
admin> set filter-list 2 mcast-ip-address = 239.255.129.120
admin> write -f
admin> dir mcast-service
      802  07/24/2005 20:12:09  bronze-service
      809  07/24/2005 20:13:19  gold-service
```

The following command displays information about the multicast service profiles:

```
admin> igmp profiles
IGMP Service Profiles

      Service Name      : gold-service
      Filter Type       : MCAST_FILTER_INCLUSIVE
      Filter List       :
                          224.255.129.120
                          224.225.129.119

      Service Name      : bronze-service
      Filter Type       : MCAST_FILTER_INCLUSIVE
      Filter List       :
                          224.255.129.119
```

### **Sample multicast address range filter**

In this example, the administrator wants to allow a user access to a set of group addresses from 234.1.1.1 to 234.1.1.5. The following commands configure the addresses in an MCAST-SERVICE profile:

```
admin> new mcast-service test_234_1
admin> set active = yes
admin> set filter-type = inclusive
admin> set filter-list 1 active = yes
admin> set filter-list 1 mcast-ip-address = 234.1.1.1
admin> set filter-list 1 group-range-count = 5
admin> write -f
```

The following commands configure the user's CONNECTION profile to use the new MCAST-SERVICE profile:

```
admin> new connection tg61
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 10.10.7.10/24
```

```
admin> set ip-options multicast-allowed = yes
admin> set ip-options multicast-service-profile = test_234_1
admin> set atm-options nailed-group = 251
admin> write -f
```

If the configuration results in an out-of-range multicast address, or if the same group address is specified in two `filter-list` subprofiles, the Stinger logs an error message.

## Configuring multicast client interfaces

A multicast client interface is an IP-enabled ATM PVC across a DSL interface to a CPE router such as a CellPipe®.

## Overview of multicast client CONNECTION settings

In addition to the ATM and IP options needed to configure a terminating PVC, the following parameters, shown with default values, are used to enable IP multicast on a client interface:

```
[in CONNECTION/"":ip-options]
multicast-allowed = no
multicast-rate-limit = 100
multicast-group-leave-delay = 0
multicast-group-leave-delay-msec = 0
multicast-max-groups = 0
fast-leave = no
multicast-server-vc = no
multiple-mcast-filter = [ "" "" "" "" "" "" "" "" "" "" "" "" "" "" "" "" "" "+]
```

| Parameter            | Setting                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| multicast-allowed    | <p>Enables or disables handling of IGMP requests and responses on the interface.</p> <p>The system does <i>not</i> forward multicast traffic on the basis of this setting.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| multicast-rate-limit | <p>The rate at which the system accepts multicast packets from clients on the interface. For example, if you set the rate to 5, the system accepts one packet every 5 seconds from multicast clients on the interface, and discards subsequent packets received within that 5-second window.</p> <p>The default value of 100 disables multicast forwarding on the interface—the Stinger IP DSLAM acts as a forwarder and handles IGMP packets, but does not accept packets from clients or forward multicast packets from the MBONE router. To enable multicast forwarding on the interface, you must set the rate to a number <i>less than</i> 100.</p> |

| Parameter                        | Setting                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| multicast-group-leave-delay      | <p>Number of seconds to delay before forwarding a Leave Group message. The sum of (multicast-group-leave-delay × 1000) plus multicast-group-leave-delay-msec is the number of milliseconds the system waits before forwarding to the MBONE router an IGMP version-2 Leave Group message it receives across a multicast client interface. With the zero default values, the system forwards the Leave Group messages immediately. For client interfaces that support multiple multicast sessions to the same group, set these parameters to a value from 10 to 20.</p> <p>When these parameters have nonzero values, the system sends back a query to make sure there are no active multicast sessions on the interface for that group, and if it receives a response before the specified, specified delay expires, it does not forward the Leave Group message.</p>                                                                                                                                                                                                                                                                                        |
| multicast-group-leave-delay-msec | <p>Number of milliseconds to add to the value of (multicast-group-leave-delay × 1000) to determine the total delay before forwarding a Leave Group message.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| multicast-max-groups             | <p>Maximum number of accessible multicast group (from 0 to 250) for this client interface. You can set this to a lower number to limit multicast traffic to the interface. This value of this parameter limits the activation of new client links, but does not affect the existing client links.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| fast-leave                       | <p>Enables or disables the fast leave feature, which bypasses the Group Leave delay. When set to <b>no</b> (the default), after a Leave Group message is received from the last group member, the system delays the Group Leave process to run a last member query, to ensure that the host is the last member. If a host sends a host membership report in response to the query, the Group Leave process is canceled and the host membership interval timer is started. The IP multicast traffic does not stop until the last host member timer expires.</p> <p>When set to <b>yes</b>, the system seizes the IP multicast traffic on an interface or sends Group Leave message to upstream router on MBONE interface as soon as the last member host leaves the multicast group.</p> <p>IGMP v2 or later is required for proper operation. This feature assumes that individual hosts are sending membership reports on receiving the general or group membership query, as is the case for well-known STBs.</p> <p>You can view active member interfaces by using the <code>igmp hosts</code> command on a LIM. See “IGMP diagnostics” on page A-4.</p> |



| Parameter             | Setting                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| multicast-server-vc   | Specifies whether the virtual circuit is a multicast server VC (yes or no). See “Overview of multicast server VC settings” on page 9-23.                                                                                                                                                                                                                                                                                                                                                                                          |
| multiple-mcast-filter | An array of 20 indexed subprofiles, to enable you to apply up to 20 mcast-service profiles per user. The profiles are applied in the order in which they are specified. Once a group address is found in an MCAST-SERVICE profile in either the include list or exclude list, it is applied. If at least one multicast filter exists and is active, groups that are not in the include list are not allowed. If no MCAST-SERVICE profiles are configured or none of the profiles are active, then all groups are allowed to join. |

### **Setting IGMP-v2 timers (local profiles only)**

The system calculates the timeout value for IGMPv1 and IGMPv2 members using the rule described in RFC 2236 *Internet Group Management Protocol, Version 2*. The rule is as follows:

Timeout-value = ( query-interval x robust-count + query-response-interval/10 )

You can control the timeout value by setting the query-interval, robust-count, or query-response-interval parameters for an IGMPv1 or IGMPv2 member.

The following parameters, shown with default values, are used to configure the timers defined in RFC 2236 on multicast client interfaces:

```
[in CONNECTION/"";ip-options:igmp-options]
robust-count = 2
query-interval = 125
query-response-interval = 100
last-member-query-interval = 10
last-member-query-count = 2
```

| Parameter      | Setting                                                                                                                                                                                                                                                                                                                  |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| robust-count   | A threshold of packet losses (from 2 to 10) up to which the multicast subsystem will remain robust. If the interface is expected to have a high rate of packet loss, increase this value. IGMP is robust to (robust-count minus 1) packet losses. It cannot be set to zero and should not be set to 1. The default is 2. |
| query-interval | Number of seconds (from 0 to 1024) between general queries. You can increase this value from its default of 125 seconds to reduce the number of IGMP queries sent on the interface.                                                                                                                                      |

| Parameter                  | Setting                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| query-response-interval    | Maximum response time in tenths of a second (from 0 to 1024) inserted into general queries. You can increase this value from its default of 10 seconds to make IGMP traffic less bursty, because host responses will be spread out over a larger interval. The number of seconds response time (this value divided by 10) must be less than the query-interval value.       |
| last-member-query-interval | Maximum response time in tenths of a second (from 0 to 1024) inserted into group-specific queries sent in response to Leave Group messages. You can reduce this value from its default of 1 second to reduce the time it takes to detect that the last member of a group has left. The response time (this value divided by 10) must be less than the query-interval value. |
| last-member-query-count    | Number of group-specific queries sent before the multicast router assumes there are no local members.                                                                                                                                                                                                                                                                       |

### **Example of using multiple multicast filters**

In this example, an administrator defines one MCAST-SERVICE profile for a basic set of 50 channels to be allowed to a subscriber, and another MCAST-SERVICE profile for each available PPV event. Then, to allow a user to see a PPV event, the administrator adds the filter for the PPV event to the user's CONNECTION profile. The existing filter for the basic set of channels remains active and unchanged in the user's profile, and multiple PPV-related filters can be applied in the same profile as appropriate.

The following set of commands defines an MCAST-SERVICE profile for a basic set of channels:

```
admin> new mcast-service
admin> set service-name = 50-channels
admin> set active = yes
admin> set filter-type = inclusive
admin> set filter-list 1 active = yes
admin> set filter-list 1 mcast-ip-address = 224.10.10.10
admin> set filter-list 1 group-range-count = 50
admin> write -f
```

The following set of commands defines an MCAST-SERVICE profile for a PPV event:

```
admin> new mcast-service
admin> set service-name = ppv-07302004
admin> set active = yes
admin> set filter-type = inclusive
admin> set filter-list 1 active = yes
```

```
admin> set filter-list 1 mcast-ip-address = 226.10.10.10
admin> set filter-list 1 group-range-count = 1
admin> write -f
```

The following set of commands defines a CONNECTION profile and applies both filters:

```
admin> new connection mpoa-adsl
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 110.110.110.52/8
admin> set ip-options netmask-remote = 255.0.0.0
admin> set ip-options multicast-allowed = yes
admin> set ip-options multicast-max-groups = 250
admin> set ip-options multiple-mcast-filter 1 = 50-channels
admin> set ip-options multiple-mcast-filter 2 = ppv-07302004
admin> set atm-options nailed-group = 604
admin> write -f
```

The following commands displays information about the MCAST-SERVICE profiles:

```
admin> igmp profiles
IGMP Service Profiles
```

```
Service Name           : 50-channels
Filter Type             : MCAST_FILTER_INCLUSIVE
Filter List( Group Range) :
224.10.10.10 - 224.10.10.59 ( Count = 50 )
```

```
Service Name           : ppv-7302004
Filter Type             : MCAST_FILTER_INCLUSIVE
Filter List( Group Range) :
226.10.10.10 - 226.10.10.10 ( Count = 1 )
```


## **Sample multicast video configuration with filters**

In this sample setup, the MBONE interface is on a Gigabit Ethernet port. A multicast router on the local IP network sends video transmissions from a multicast video server. In the sample configurations for the setup shown in Figure 9-5, two types of filters are used:

- MCAST-SERVICE profiles, to filter multicast groups and related services
- FILTER profiles, to filter out generic IP data

By introducing generic IP data filters, you can cause the system's IP interfaces to handle only multicast video services and block end users from accessing other services in the network.

Figure 9-5. DSL video application with a local MBONE interface



## Configuring the local MBONE interface

The following commands enable the MBONE interface on the Gigabit Ethernet port:

```
admin> read ip-global
admin> set multicast-forwarding = yes
admin> set mbone-lan-interface 1 = { { 1 8 2 } 0 }
admin> write -f
```

## Configuring multicast client PVCs

The profiles in this section apply MCAST-SERVICE profiles to filter multicast groups and their related services. For background information, see “Managing multicast group memberships” on page 9-8.

The following commands configure a PVC for CPE-1 in Figure 9-5:

```
admin> new connection cpe-1
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 2.2.2.2/29
admin> set ip-options multicast-allowed = yes
admin> set ip-options multicast-rate-limit = 20
admin> set ip-options multicast-service-profile = bronze-service
admin> set ip-options multicast-max-groups = 1
admin> set atm-options vci = 100
admin> set atm-options nailed-group = 51
admin> write -f
```

The next set of commands configures the CPE-2 client interface in Figure 9-5:

```
admin> new connection cpe-2
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 3.3.3.3/29
```

```
admin> set ip-options multicast-allowed = yes
admin> set ip-options multicast-rate-limit = 20
admin> set ip-options multicast-service-profile = gold-service
admin> set ip-options multicast-max-groups = 2
admin> set atm-options vci = 100
admin> set atm-options nailed-group = 52
admin> write -f
```

## Applying a filter that restricts the GigE interface to video traffic only

In this example, a filter is applied to the GigE interface to prevent the interface from handling traffic other than video data and related control messages. This filter has the advantage of being easy to define and apply, but it restricts the GigE interface from being used for other applications.

For background information about filters, see Chapter 11, “Filter Configuration.”

The following commands create a new FILTER profile named `mcast-only`, and specify an input filter that forwards inbound multicast data traffic (239.100.100.0):

```
admin> new filter mcast-only
admin> set input-filters 1 valid-entry = yes
admin> set input-filters 1 forward = yes
admin> set input-filters 1 Type = ip-filter
admin> set input-filters 1 ip-filter source-address-mask = 255.255.255.0
admin> set input-filters 1 ip-filter source-address = 239.100.100.0
```

The following commands specify an output filter that allows client access to the video server at IP address 5.5.5.5 (Figure 9-5). Clients might require access to the server, for example, to download boot information for a set-top box.

```
admin> set output-filters 1 valid-entry = yes
admin> set output-filters 1 forward = yes
admin> set output-filters 1 Type = ip-filter
admin> set output-filters 1 ip-filter dest-address-mask = 255.255.255.255
admin> set output-filters 1 ip-filter dest-address = 5.5.5.5
```

Because the default setting of the `forward` parameter is `no`, the next set of commands explicitly drops all other output traffic:

```
admin> set output-filters 2 valid-entry = yes
admin> set output-filters 2 Type = ip-filter
admin> write -f
```

The following commands apply the `mcast-only` filter to the Gigabit Ethernet interface:

```
admin> read ethernet { 1 8 2 }
admin> set filter-name = mcast-only
admin> write -f
```

## An alternative filter to restrict each client interface

If you must use a Gigabit Ethernet interface for other applications as well as multicast video, you cannot restrict the type of traffic allowed on the interface. In that case, you can define filters for individual multicast client interfaces, to restrict the those interfaces from handling traffic other than video data and related control messages. Individual users on the client interface will be unable to access other services through the Stinger IP DSLAM.



**Note** The Stinger IP DSLAM does not currently support filters specified in RADIUS profiles.

This sample filter is specific to the multicast client network connected to CPE-1 in Figure 9-5 (page 9-16), with the CPE IP address 2.2.2.2/29.

The following commands create a new FILTER profile named `conn-input-filter`, and specify an input filter that forwards inbound traffic to the video server at IP address 5.5.5.5 (Figure 9-5). Clients might require access to the server, for example, to download boot information for a set-top box.

```
admin> new filter conn-input-filter
admin> set input-filters 1 valid-entry = yes
admin> set input-filters 1 forward = yes
admin> set input-filters 1 Type = ip-filter
admin> set input-filters 1 ip-filter dest-address-mask = 255.255.255.255
admin> set input-filters 1 ip-filter dest-address = 5.5.5.5
```

The next set of commands specifies an input filter that allows only inbound traffic that uses the IGMP protocol (protocol number 2) from a client on the subnet 2.2.2.x.

```
admin> set input-filters 2 valid-entry = yes
admin> set input-filters 2 forward = yes
admin> set input-filters 2 Type = ip-filter
admin> set input-filters 2 ip-filter protocol = 2
admin> set input-filters 2 ip-filter source-address-mask = 255.255.255.0
admin> set input-filters 2 ip-filter source-address = 2.2.2.0
```

Because the default setting of the `forward` parameter is `no`, the next set of commands explicitly drops all other inbound traffic from the client subnet:

```
admin> set input-filters 3 valid-entry = yes
admin> set input-filters 3 Type = ip-filter
admin> write -f
```

For the CPE-1 CONNECTION profile definition, see “Configuring multicast client PVCs” on page 9-16. The following commands apply this filter to the individual client interface for CPE-1:


```
admin> read connection cpe-1
admin> set session-options data-filter = conn-input-filter
admin> write -f
```

## Sample multicast video configuration with a remote MBONE interface

In the sample setup shown in Figure 9-6, the MBONE interface is configured in Stinger-2, and the multicast client interface is configured in Stinger-1.

The connection between the two Stinger units is an ATM PVC. It can use any ATM medium, and does not use IP processing. In this example, Stinger-2 sends the IPTV multicast data stream across the ATM cloud as a cell stream. Stinger-1 forwards the ATM stream to the multicast CPE router on the basis of an ATM circuit configuration.

Figure 9-6. IPTV video sample configuration



Stinger-1 requires an ATM circuit profile between the LIM interface to the CPE router, and the trunk interface to Stinger-2. For related information, see “Distributed multicast with HB LIMs (IP2100 only)” on page 9-25.

The following command on Stinger-1 displays the nailed-group number of the ADSL interface (LIM slot 4, port 5) to the CPE router in Figure 9-6:

```
admin> which -n { 1 4 5 }  
Nailed group corresponding to port { shelf-1 slot-4 5 } is 155
```

The following command on Stinger-1 displays the nailed-group number of the OC3-ATM interface (trunk slot 17, port 2) that connects to the ATM core network in Figure 9-6:

```
admin> which -n { 1 17 2 }  
Nailed group corresponding to port { shelf-1 trunk-module-1 2 } is 802
```

The following set of commands on Stinger-1 configures an ATM circuit between the two interfaces:

```
admin> new connection mcast-client-pvc  
admin> set active = yes  
admin> set encapsulation-protocol = atm-circuit  
admin> set ip-options ip-routing-enabled = no  
admin> set atm-options nailed-group = 155  
admin> set atm-connect-options vci = 100  
admin> set atm-connect-options nailed-group = 802  
admin> write -f
```

Following is a comparable RADIUS profile:

```
permconn-st-2 Password = "pwd"
  Service-Type = Outbound,
  Framed-Protocol = ATM-CIR,
  User-Name = "mcast-client-pvc",
  Ascend-ATM-Group = 155,
  Ascend-Route-IP = Route-IP-No,
  Ascend-ATM-Vpi = 0,
  Ascend-ATM-Vci = 35,
  Ascend-ATM-Connect-Vpi = 0,
  Ascend-ATM-Connect-Vci = 100,
  Ascend-ATM-Connect-Group = 802
```

With this CONNECTION or RADIUS profile, ATM cells received by Stinger-1 from the CPE router are switched to the unit's trunk interface and transmitted across the ATM cloud. On Stinger-2, the PVC terminates and is packetized for transmission on the IP network.

The following command on Stinger-2 displays the nailed-group number of a DS3-ATM interface (trunk slot 17, port 1) that connects to the ATM core network in Figure 9-6:

```
admin> which -n { 1 17 1 }
Nailed group corresponding to port { shelf-1 trunk-module-1 1 } is 801
```

The following set of commands on Stinger-2 terminates the PVC on a trunk interface with VCI 100, and specifies the required IP address for Stinger-1.

```
admin> new connection term-18-100
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 3.3.3.3/29
admin> set atm-options vci = 100
admin> set atm-options nailed-group = 801
admin> write -f
```

Following is a comparable RADIUS profile:

```
permconn-st-1 Password = "pwd"
  Service-Type = Outbound,
  Framed-Protocol = ATM-1483,
  User-Name = "term-18-100",
  Framed-IP-Address = 3.3.3.3,
  Framed-IP-Netmask = 255.255.255.252,
  Ascend-ATM-Group = 801,
  Ascend-Route-IP = Route-IP-Yes,
  Ascend-ATM-Vpi = 0,
  Ascend-ATM-Vci = 100
```

After completing the configuration, starting the IPTV client software on the multicast client interface should create the 239.100.100.4 multicast group on Stinger-2.



The following command on Stinger-2 checks that the group exists:

```
admin> igmp groups
Group Address      Members    Expire time    Counts
230.0.0.9          14         00:00:31       0 :: 0 S2
                  *(Mbone)    0 :: 0 S2
```

The following command displays client interfaces (interface 14 representing the remote client interface):

```
admin> igmp clients
IGMP Clients
Client            Version    RecvCount      CLU    ALU
1(Mbone)          2         0              0      0
14                2         0              0      0
```

## ATM QoS when both multicast and unicast clients are supported

When the system is processing both multicast and unicast data streams and one or both of the following conditions exists, some data discards will be observed on the unicast connections.

- When multicast data is sent at a relatively higher rate
- When many multicast clients have joined while unicast data is being sent

For example, suppose a LIM in the Stinger system has the subscriber connections and data rates shown in Table 9-2:

Table 9-2. Sample multicast and unicast client requirements

| Type of traffic Data rate |                     | Number of clients or connections | Total data rate |
|---------------------------|---------------------|----------------------------------|-----------------|
| Multicast                 | 6.51Mbps/client     | 22                               | 143.22Mbps      |
| Multicast                 | 5.80Mbps/client     | 23                               | 134Mbps         |
| Unicast                   | 5.80Mbps/connection | 47                               | 274Mbps         |
|                           |                     |                                  | 551Mbps         |

With the client requirements shown in Table 9-2, you must use ATM QoS to enable the system to shape the unicast traffic at the desired rate while multicast data is also being transmitted. Following are the ATM QoS guidelines for unicast connections:

- 1 Configure the sustained cell rate (SCR) approximately 5% higher than the desired data send rate on the connection.
- 2 Configure the peak cell rate (PCR) at twice the SCR value.
- 3 Set the maximum burst size to 30.

For example, the following ATM-QOS profile follows these guidelines. Note that the PCR (peak-rate-kbits-per-sec) is twice the rate of the SCR, and the SCR (sustainable-rate-kbits-per-sec) is the desired rate times 1.05 ( $5.8 * 1.05 = 6000$ ). The max-burst-size setting is 30.

```
admin> get atm-qos stng-ucast-ds
[in ATM-QOS/stng-ucast-ds]
contract-name* = stng-ucast-ds
traffic-descriptor-index = 32
```

```
traffic-descriptor-type = noclpscr
atm-service-category = real-time-vbr
peak-rate-kbits-per-sec = 12000
peak-cell-rate-cells-per-sec = 28301
sustainable-rate-kbits-per-sec = 6000
sustainable-cell-rate-cells-per-sec = 14150
minimum-rate-kbits-per-sec = 0
minimum-cell-rate-cells-per-sec = 0
ignore-cell-delay-variation-tolerance = no
cell-delay-variation-tolerance = 20
ignore-max-burst-size = no
max-burst-size = 30
aal-type = aal-5
early-packet-discard = yes
partial-packet-discard = no
tag-or-discard = discard
external-change = no
sub-channel = 1
```


This ATM-QOS profile can be applied to unicast connections to enable the system to shape unicast traffic at the desired rate at the same time that multicast data is being transmitted.

## Multicast server virtual circuits

With this software version, you can configure up to 100 trunk PVCs as *multicast server virtual circuits (VCs)*. A multicast server VC is a trunk terminating PVC that carries the multicast data stream of one multicast group.

The administrator of an upstream router maps a PVC to a particular multicast group address and ensures that only one multicast group is mapped to one VC, as shown in Figure 9-7. The Stinger system does not need to know which group address is carried on each VC to establish the connection.

Figure 9-7. One VC per multicast group for incoming multicast data streams



When the multicast server VC has been established in the Stinger IP DSLAM, it is registered as an MBONE interface. The VC acts as an MBONE interface for IGMP joins coming from the DSL side. However, IGMP control messages are not exchanged across the multicast server VC. The system simply forwards multicast traffic received

on the multicast server VC to all DSL clients who have joined the associated multicast group.



**Note** When `multiple-multicast-server-vc` is configured in the IP-GLOBAL profile, you cannot also configure MBONE or PIM.

## Overview of multicast server VC settings

Following are the parameters, shown with default settings, required for configuring multicast server VCs:

```
[in IP-GLOBAL]
multicast-forwarding = no
multiple-multicast-server-vc = no


[in CONNECTION/"" :ip-options]
multicast-allowed = no
multicast-server-vc = no
```

| Parameter                    | Setting                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| multicast-forwarding         | Enables or disables multicast forwarding. When you change the value to <b>yes</b> and write the profile, the multicast subsystem reads the values in the IP-GLOBAL profile and initiates the forwarding function. Set to <b>yes</b> for the multicast server VC feature.                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| multiple-multicast-server-vc | Enables or disables the multicast server VC feature systemwide. With the default <b>no</b> setting, the feature is disabled. Set it to <b>yes</b> to enable the feature.<br><br>When this parameter is set to <b>yes</b> , you can configure up to 100 trunk PVCs Stinger to receive multicast data from one multicast group.<br><br>When this parameter is set to <b>yes</b> , you cannot configure MBONE or PIM.                                                                                                                                                                                                                                                                                                            |
| multicast-allowed            | Enables or disables the handling of IGMP requests and responses on trunk MBONE interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| multicast-server-vc          | Specifies whether the virtual circuit is a multicast server VC ( <b>yes</b> or <b>no</b> ). A multicast server VC is a trunk terminating PVC that carries the multicast data stream of one multicast group.<br><br>It acts as an MBONE for IGMP joins coming from the DSL side, and the system forwards multicast traffic from the multicast server VC to DSL clients who have joined the group. However, IGMP control messages are not exchanged on the interface.<br><br>When this parameter is set to <b>yes</b> , the <code>encapsulation-protocol</code> parameter must be set to <code>atm</code> .<br><br>The <code>multiple-multicast-server-vc</code> parameter in the IP-GLOBAL profile must be set to <b>yes</b> . |

## Sample configuration of multicast server VCs

In the sample setup shown in Figure 9-8, STREAM-1, STREAM-2, STREAM-3 are trunk terminated ATM PVCs configured as multicast server VCs, and user-1, user-2, and user-3 are DSL user connections.

Figure 9-8. Multicast server VCs on a trunk interface



In this example, the Stinger system is receiving the multicast data stream of one multicast group on each of the three trunk-terminating PVCs. An upstream router (not shown in this example) is responsible for streaming multicast traffic on each VC.

### Enabling the multicast server VC feature

The following commands on the Stinger system enable the multicast server VC feature systemwide:

```
admin> read ip-global
admin> set multicast-forwarding = yes
admin> set multiple-multicast-server-vc = yes
admin> write -f
```

### Configuring the multicast server VCs

The following commands create the STREAM-1, STREAM-2, STREAM-3 trunk-terminated ATM PVCs shown in Figure 9-8. The `multicast-server-vc` setting must be enabled in the VC CONNECTION profiles, and `encapsulation-protocol` must be set to `atm`.

```
admin> new connection STREAM-1
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 20.20.20.1/8
admin> set ip-options multicast-allowed = yes
admin> set ip-options multicast-server-vc = yes
admin> set atm-options vci = 40
admin> set atm-options nailed-group = 802
admin> write -f
admin> new connection STREAM-2
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 20.20.20.2/8
```

```
admin> set ip-options multicast-allowed = yes
admin> set ip-options multicast-server-vc = yes
admin> set atm-options vci = 41
admin> set atm-options nailed-group = 802
admin> write -f
admin> new connection STREAM-3
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 20.20.20.3/8
admin> set ip-options multicast-allowed = yes
admin> set ip-options multicast-server-vc = yes
admin> set atm-options vci = 43
admin> set atm-options nailed-group = 802
admin> write -f
```

### Configuring the terminating connections for DSL users

Following is a sample terminating CONNECTION profile for the connection labeled user-1 in Figure 9-8:

```
admin> new connection user-1
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 122.122.122.122/8
admin> set ip-options local-address = 111.111.111.111/8
admin> set ip-options multicast-allowed = yes
admin> set atm-options vci = 59
admin> set atm-options nailed-group = 505
admin> write -f
```

## ***Distributed multicast with HB LIMs (IP2100 only)***

Stinger systems with an IP2100 control module and high-bandwidth (HB) Annex A LIMs with ADSL2+ capability (STGR-LIM-AP-72-HB or STGR-LIM-AP-48-HB) support a distributed multicast mechanism.

### ***Benefits of distributed multicast***

Distributed multicast provides the following benefits:

- With distributed multicast, much less backplane bandwidth is required for multicast video streams.

With centralized (non-distributed) multicast, the system uses backplane bandwidth for each multicast client. With large numbers of clients sourcing video at the same time, the throughput requirement could exceed the maximum backplane bandwidth. For example, the video throughput required for MPEG2 is approximately 3Mbps. If 3\*72 (216) multicast clients are sourcing MPEG2, the

## IP Multicast Configuration

### Distributed multicast with HB LIMs (IP2100 only)

resulting bandwidth requirement (648Mbps) would exceed the backplane bandwidth of 622Mbps.

With distributed multicast, the system uses backplane bandwidth for each multicast *channel*. So up to 622/3 (207) simultaneous video channels can be transmitted to a LIM regardless the number of clients watching a channel.

- With distributed multicast, much less demand is placed on the control module's network processor resources.

With centralized multicast, the control module manages all client interfaces. With distributed multicast, the control module manages only 14 clients—the 14 slots. The load placed on the control module's network processor is dramatically decreased.

- With distributed multicast, many more multicast clients can be serviced.

With centralized multicast, up to 1017 multicast clients can be supported at the same time. With distributed multicast, up to 216 multicast clients (3\*72) can be supported on each LIM, so up to 3024 (14\*216) video clients can be supported.

## Conceptual overview

The distributed multicast mechanism allows the following types of module to replicate multicast streams:

- IP2100 control module (CM)
- High-bandwidth (HB) LIMs


Each of these multicast-capable modules acts as a multicast replication point and operates as a standalone multicast router that manages its own multicast interfaces.



**Note** For bridged connections, only the control module performs multicast stream replication.

Modules acting as multicast routers are interconnected through special interfaces called Fast Forwarding (FFW) interfaces. The FFW interfaces are not user configurable, and are unnumbered (they do not have an IP address). Figure 9-9 shows the interfaces for two multicast-capable modules in a Stinger system, the IP2100 CM and an HB LIM.

Figure 9-9. FFW interfaces on multicast-capable modules



From a multicast point of view, the CM does not manage all interfaces in the system.

- The multicast-capable CM manages its GigE ports and its own FFW interfaces.
- The multicast-capable HB LIM manages its own WAN interfaces and its FFW interface.

## MBONE and multicast client interactions on FFW interfaces

Table 9-3 shows the MBONE and multicast client behavior of the FFW interfaces:

Table 9-3. FFW interfaces registered as MBONE or multicast client

| Module | FFW interfaces                                                                                             | FFW multicast type |
|--------|------------------------------------------------------------------------------------------------------------|--------------------|
| IP2100 | One FFW interface to each slot in the system. The internal port is the physical port linked to the module. | multicast-client   |
| HB LIM | One FFW interface to the CM                                                                                | multicast-mbone    |

### What happens when a CPE joins a multicast group


When an IGMP Join message is received, each multicast router performs the following steps:

- 1 If the multicast group does not exist on the module, create it. Otherwise, skip to step 4.
- 2 Register the MBONE interfaces to this group.
- 3 Proxy the IGMP Join on the MBONE interfaces.
- 4 Connect the client to the multicast group.

For example, in Figure 9-10, two CPEs join multicast group 224.1.1.1.

When the system first receives an IGMP Join message from CPE-1, both the HB LIM and the IP2100 CM perform all four steps listed above. Once the group has been created, additional Join messages cause the system to simply connect the client to the multicast group at the HB LIM.

Figure 9-10. Two IGMP Join requests for the same multicast group



If one of the CPEs in Figure 9-10 issues another IGMP Join message for a different multicast group (for example, a Join message for 224.2.2.2), the process is repeated for the new group.





- Each of the HB LIM's DSL lines supports a maximum of four terminated PVCs or up to eight switched ATM PVCs.
- OAM support should be limited to no more than 400 connections on a LIM if provisioning the maximum number of eight switched ATM PVCs per port while using unique VPI values for each of the connections on each port.

## Changes in packet-flow prioritization with distributed multicast

For all LIMs other than the HB LIM, the control module handles both unicast and multicast traffic and their associated QoS.

For the HB LIM, the control module still handles the output of unicast traffic and its associated QoS, but the downstream output of multicast traffic is managed by the HB LIM. This has introduced changes in how downstream traffic on the HB LIM is prioritized.

On the HB LIM, with a **packet-flows** profile that prioritizes a mix of unicast and multicast traffic flows, downstream traffic is prioritized following these rules:

- If the highest priority flow in the **packet-flows** profile is unicast traffic, the HB LIM gives all unicast traffic a higher priority than all multicast traffic.
- If the highest priority flow in the **packet-flows** profile is multicast traffic, the HB LIM gives all multicast traffic a higher priority than all unicast traffic.
- If both multicast and unicast traffic share the highest priority in the **packet-flows** profile, the HB LIM gives all multicast traffic a higher priority than all unicast traffic.

Also note that the QoS marking of multicast traffic is not supported in the downstream direction on the HB LIM.

Following are some examples of how the HB LIM prioritizes traffic for **packet-flows** profiles:

*Table 9-4. Packet flow with multicast traffic as the highest priority flow*

| Flow #        | Description of classifier          | Priority   | Result                                   |
|---------------|------------------------------------|------------|------------------------------------------|
| <b>flow 1</b> | Multicast traffic classifier       | priority 7 | These priorities are strictly respected. |
| <b>flow 2</b> | Unicast traffic classifier         | priority 6 |                                          |
| <b>flow 3</b> | Another unicast traffic classifier | priority 5 |                                          |

*Table 9-5. Packet flow with unicast traffic as the highest priority flow*

| Flow #        | Description of classifier        | Priority   | Result                                   |
|---------------|----------------------------------|------------|------------------------------------------|
| <b>flow 1</b> | Unicast traffic classifier       | priority 7 | These priorities are strictly respected. |
| <b>flow 2</b> | Unicast traffic classifier       | priority 6 |                                          |
| <b>flow 3</b> | Default rule (unicast/multicast) | priority 5 |                                          |

*Table 9-6. Packet flow with interleaving priorities (unicast/multicast)*

| Flow #        | Description of classifier        | Priority   | Result                                                              |
|---------------|----------------------------------|------------|---------------------------------------------------------------------|
| <b>flow 1</b> | Unicast traffic classifier       | priority 7 | All unicast traffic will be higher priority than multicast traffic. |
| <b>flow 2</b> | Multicast traffic classifier     | priority 6 |                                                                     |
| <b>flow 3</b> | Default rule (unicast/multicast) | priority 5 |                                                                     |

*Table 9-7. Packet flow with interleaving priorities (multicast/unicast)*

| Flow #        | Description of classifier        | Priority   | Result                                                              |
|---------------|----------------------------------|------------|---------------------------------------------------------------------|
| <b>flow 1</b> | Multicast traffic classifier     | priority 7 | All multicast traffic will be higher priority than unicast traffic. |
| <b>flow 2</b> | Unicast traffic classifier       | priority 7 |                                                                     |
| <b>flow 3</b> | Default rule (unicast/multicast) | priority 6 |                                                                     |

## Scheduling priority of the HB LIM FFW interface

By default, the HB LIM FFW interface is assigned a CBR service category for delivering multicast traffic to the LIM. This means that if a BIR connection that carries both multicast video and unicast video is assigned a VBR service category (realtime or non-realtime), the multicast traffic is still scheduled at CBR (highest) priority while the unicast traffic is scheduled at VBR priority. As a result, the unicast traffic is not serviced fast enough and data discards can occur.

To solve an imbalance in scheduling priorities for multicast and unicast video, you must reset the FFW interface scheduling priority so both types of traffic are scheduled at the same priority level. To do so, use the following parameter, shown with its default value:

```
[in SLOT-STATIC-CONFIG/{ shelf-1 slot-4 0 }]
ffw-scheduling-priority = scheduler-priority-0
```

| Parameter               | Setting                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ffw-scheduling-priority | Specifies the FFW interface scheduling priority for multicast video traffic.<br><br>Because the FFW interface is set up during the LIM POST stage, a change in this parameter's value requires a slot bounce (by using slot -b) for the change to take effect. Following are valid values:<br><div> <div>scheduler-priority-0</div> <div>The highest priority (CBR). This is the default setting.</div> <div>scheduler-priority-1</div> <div>Priority 1 (rt-VBR).</div> <div>scheduler-priority-2</div> <div>Priority 2 (nrt-VBR).</div> <div>scheduler-priority-3</div> <div>Priority 3 (UBR). The lowest priority.</div> </div> |

## ATM QoS for unicast traffic on the HB LIM

The HB LIM performs multicast stream replication. However, the same considerations apply as when the IP control module is performing the replication. For details about how to ensure accurate shaping of unicast traffic when both multicast and unicast streams are being processed, see “ATM QoS when both multicast and unicast clients are supported” on page 9-21.

## Statistics and diagnostics for the multicast-capable HB LIM

The system provides several commands for obtaining statistics and debug-level diagnostics related to multicast operations on the HB LIM. To use these commands, you must first open an administrative session to the module. For example, the following commands open a session to a module in slot 15, and then exit the session:

```
admin> open 1 15
ep-hs-gs-48-1/15> quit
```

The following commands provide information about distributed multicast on a multicast-capable HB LIM:

- The stats commands with the **ffwBridge** option now display traffic, error, and rate information related to distributed multicast on the HB LIMs. For details, see the *Stinger Reference*.
- The stats traffic mcast command can now use the -g <group> or the -i <if-num> option to display multicast traffic statistics about a specific multicast group or interface on the HB LIMs. For details, see the *Stinger Reference*.
- The debug-level info command on the HB LIMs displays information about the internal AAL5 bridge, and about multicast traffic in proxy (routing) or snooping (bridging) mode. For details, see “Multicast-capable high-bandwidth LIM diagnostics” on page A-27.

## Administrative tools for IGMP operations

The system supports the **igmp** commands for administrative information about IGMP multicast operations. For examples, see “IGMP diagnostics” on page A-4.

## CAC for multicast channels

CAC for Multicast Channels lets you manage bandwidth of the Gigabit Ethernet interfaces present on the Stinger control module. If you consider each join received on Stinger as a connection request, then multicast CAC can be considered analogous to ATM CAC. You can assign a bandwidth for each multicast channel.

- Provides a way to allocate a fixed amount of bandwidth exclusively for multicast traffic. Reserve the bandwidth either during provisioning or when a join is received for this channel. This is subject to multicast CAC constraints.
- Support for Multicast CAC is available only on the Gigabit interfaces present on the Stinger Control Module.
- This feature is supported on the Stinger IP2000 and IP2100 Control Modules.



**Note** Bandwidth reservation is done only on the control plane. This means, no bandwidth is actually reserved for multicast traffic in the network processor. The network processor is unaware of any CAC configuration.

## MULTICAST-CAC

The `multicast-cac` profile configures the CAC rules. The CAC rules apply to the first join for each multicast group that the system receives. If the system accepts the first join for a configured group, then it accepts all joins for the same group.

The system currently supports only a single `multicast-cac` profile. This profile applies to both IGMP proxy and snooping configurations.

The following is a listing of the profile with its default settings:

```
admin> read multicast-cac
MULTICAST-CAC read
admin> list
[in MULTICAST-CAC]
name = default
bw-allocated = 0
multicast-cac-list = [ { 0.0.0.0 no 0 } { 0.0.0.0 no 0 } { 0.0.0.0 no 0 } {
0.0+
```

The following table provides the details of the `multicast-cac` parameters:

| Parameter | Description                                                                                                                                                                                                                           |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| name      | The multicast CAC profile name.<br>You can enter up to 31 characters.<br>This is set to <code>default</code> .<br>This field is not used because support is currently available for only a single <code>multicast-cac</code> profile. |

| Parameter    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| bw-allocated | <p>Total bandwidth allocated to this multicast group bundle.</p> <p>This parameter is similar to the <code>multicast-cac-max-bw</code> in the <code>SYSTEM</code> profile.</p> <p>The value for this parameter should <i>always</i> match the value of <code>multicast-cac-max-bw</code> in the <code>SYSTEM</code> profile.</p> <p>The range of the parameter <code>bw-allocated</code> is based on the CM.</p> <p>On an IP2000 control module, range = 1 to 100000</p> <p>On and IP2100 control module, range = 1 to 200000</p> <p><b>Note</b> This parameter provides for future compatibility. When there are multiple <code>multicast-cac</code> profiles, each profile can have its maximum bandwidth allocation done through this field. The sum of bandwidths of this field for each of the profiles would be less than or equal to the <code>multicast-cac-max-bw</code> value in the <code>SYSTEM</code> profile.</p> |

| Parameter          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| multicast-cac-list | <p>Multicast group address CAC list.</p> <p>The entries contain the details of each multicast channel that is subjected to CAC constraints. If <b>multicast-cac</b> is enabled on the <b>SYSTEM</b> profile, then join for any group that is not configured in this list is dropped.</p> <p>The details of each multicast group address contains the following entries:</p> <p><b>mcast-ip-address:</b> The multicast group IP address or netmask field. For example, 235.1.1.1.</p> <p>Use only a valid class D IP address. To un-set an entry, assign 0.0.0.0 as the IP address. When you enter 0.0.0.0 as a value for the IP address, the value will not be considered for any CAC calculations even if the priority/bw-req fields have some valid values in them.</p> <p><b>bw-req:</b> The amount of bandwidth to be reserved for this multicast IP address (in kbps). Range = 1 to 200000.</p> <p><b>priority:</b> Indicates if this multicast-group is a prioritized one.</p> <ul style="list-style-type: none"><li>■ The bandwidths of priority channels are reserved during provisioning.</li><li>■ The sum of bandwidths of all the priority channels is validated against the maximum system allocated bandwidth as set in <b>multicast-cac-max-bw</b> parameter of the <b>SYSTEM</b> profile.</li><li>■ If the resulting sum is valid, then these multicast groups will always have their bandwidths reserved for use. However, if the resulting sum exceeds the <b>multicast-cac-max-bw</b>, then an error message is displayed and the user cannot write the profile.</li><li>■ The non-priority channels are allocated bandwidth dynamically, based on the available bandwidth.</li></ul> <p><b>Note</b> You cannot modify the values for this parameter if one or more joins exist for it on the system. If you attempt to modify this parameter when one or more join for it exist on the system, and the profile write fails and an error message similar to the following displays:</p> <p>error: Multicast group 235.1.0.1 is already joined on the system. Its parameters cannot be changed.</p> |

## Multicast CAC specific parameters in the SYSTEM profile

Three parameters, `multicast-cac`, `multicast-cac-max-bw`, and `multicast-cac-used-bw` in the SYSTEM profile are used to configure multicast CAC.

Following is a listing of the SYSTEM profile with the default settings. Note the three additional parameters added to configure multicast CAC.

```
SYSTEM read
admin> list
[in SYSTEM]
name = ""
...
...
multicast-cac = disable
multicast-cac-max-bw = 0
multicast-cac-used-bw =
```

## Future Provisions

| Parameter             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| multicast-cac         | <p>Enables or disables multicast CAC on the system.</p> <p><b>Warning</b> Enabling this parameter removes all multicast clients that have joined on the system. These multicast clients need to resend their joins to be able to receive traffic, subject to multicast CAC. This resets the multicast CAC statistics also.</p> <p>Valid values are disable and enable.</p> <p><b>Note</b> Set this parameter <i>before</i> enabling any multicast configuration on the system.</p> |
| multicast-cac-used-bw | <p>This is a read-only parameter. It displays the multicast bandwidth in use. It is the sum of bandwidths of all configured priority and non-priority groups that have joined on the system.</p> <p>The maximum value of <code>multicast-cac-used-bw</code> is same as the value for <code>multicast-cac-maxbw</code> for that CM.</p> <p>If <code>multicast-cac</code> is disabled on the system, then this value is set to 0.</p>                                                |
| multicast-cac-max-bw  | <p>Sets the maximum bandwidth (in kbps) that you can allocate for multicast traffic on the system.</p> <p>The maximum value of this field is based on the number of GigaBit Ethernet interfaces that the Control Module supports.</p> <ul style="list-style-type: none"> <li>■ For the IP2000, this maximum value can be set up to 100000 kbps because it has only a single GigE interface.</li> <li>■ For the IP2100 CM this can be set to a maximum of 200000 Kbps.</li> </ul>   |

New parameters have been added in the VLAN-ETHERNET and IP-INTERFACE profiles. The new parameters are not used currently. They have introduced for future compatibility.

- In the VLAN-ETHERNET: bridging-options profile, the multicast-cac-profile parameter has been added. The following is a listing of the profile with its default settings.

```
admin> new vlan-ethernet {{1 8 2} 100}
VLAN-ETHERNET/{ { shelf-1 first-control-module 2 } 100 }
read
admin> list
[in VLAN-ETHERNET/{ { shelf-1 first-control-module 2 } 100 } (new)]
interface-address* = { { shelf-1 first-control-module 2 } 100 }
vlan-id = 100
enabled = no
filter-name = ""
pppoe-options = { no no "" }
```



```
bridging-options = { 0 no no vlan-circuit 0 0 "" 0 none 0 default }  
scheduling = { 0 0 }
```

```
admin> list bridging-options  
[in VLAN-ETHERNET/{ { shelf-1 first-control-module 2 } 100 }:bridging-  
options  
(+  
bridging-group = 0  
bridge = no  
dial-on-broadcast = no  
bridge-type = vlan-circuit  
vlan-stack-user-vlan-id = 0  
qos-interface-group = 0  
packet-flows = ""  
mac-address-learning-limit = 0  
packet-type = none  
bras-id = 0  
multicast-cac-profile = default
```

- In the IP-INTERFACE profile, the multicast-cac-profile parameter is introduced. The following is a listing of the profile with its default settings.

```
admin> new ip-in {{1 8 2} 200}  
IP-INTERFACE/{ { shelf-1 first-control-module 2 } 200 } read  
admin> list  
[in IP-INTERFACE/{ { shelf-1 first-control-module 2 } 200 } (new)]  
interface-address* = { { shelf-1 first-control-module 2 } 200 }  
ip-address = 0.0.0.0/0  
netmask = 0.0.0.0  
proxy-mode = Off  
rip-mode = routing-off  
route-filter = ""  
rip2-use-multicast = yes  
ospf = { no 0.0.0.0 normal 10 40 5 simple ascend0 0 1 16777215 type-1  
c0:00:00:0:  
multicast-allowed = no  
igmp-options = { 2 125 100 10 2 v2 }  
multicast-rate-limit = 100  
multicast-group-leave-delay = 0  
multicast-group-leave-delay-msec = 0  
multicast-service-profile = ""  
multicast-max-groups = 0  
multicast-cac-profile = default  
directed-broadcast-allowed = yes  
vrouter = ""  
management-only-interface = no  
vlan-enabled = no  
vlan-id = 0  
pim-options = { no 30 105 yes 1 60 210 yes 5000 2500 }  
fast-leave = no  
bootp-relay-enable = yes  
qos-interface-group = 0  
packet-flows = ""  
scheduling = { 0 0 }
```

## Configuring multicast CAC on Stinger CM

To reserve 800 Mbps bandwidth for multicast traffic flow into the system from the GigE interface, proceed as follows:

- 1 In the system profile, enable the `multicast-cac` parameter.
- 2 Allocate the amount of bandwidth to be reserved in `multicast-cac-max-bw` parameter as in the following example:

```
admin> read system
admin> set multicast-cac = enable
admin> set multicast-cac-max-bw = 800000
admin> write -f
```

Consider, you need to allocate the available system bandwidth for multicast traffic as follows:

| Multicast group | Priority | Bandwidth required to be reserved (in Kbps) |
|-----------------|----------|---------------------------------------------|
| 234.1.2.3       | Yes      | 200000                                      |
| 234.1.2.4       | No       | 300000                                      |
| 234.1.2.5       | No       | 300000                                      |
| 234.1.2.6       | Yes      | 200000                                      |

**Note** It is possible to oversubscribe the non-priority channels because their requested bandwidths are not reserved at provisioning time. In the above example we are over-subscribing the bandwidth available for non-priority channels by 100 Mbps. Bandwidth allocation to non-priority channels is done strictly on a first-come-first-serve basis.

- 3 Configure the multicast CAC profile based on the following example.

```
admin> new MULTICAST-CAC
admin> set bw-allocated = 800
admin> set multicast-cac-list 1 mcast-ip-address = 234.1.2.3
admin> set multicast-cac-list 1 priority = yes
admin> set multicast-cac-list 1 bw-req = 200000
admin> set multicast-cac-list 2 mcast-ip-address = 234.1.2.4
admin> set multicast-cac-list 2 bw-req = 300000
admin> set multicast-cac-list 3 mcast-ip-address = 234.1.2.5
admin> set multicast-cac-list 3 bw-req = 300000
admin> set multicast-cac-list 4 mcast-ip-address = 234.1.2.6
admin> set multicast-cac-list 4 priority = yes
admin> set multicast-cac-list 4 bw-req = 200000
admin> write -f
```

When IGMP Join is dropped because of the multicast CAC, the system displays a message similar to the following:

LOG notice, Shelf 1, Controller-1, Time: 12:09:33--

Multicast client 239.255.255.250 dropped due to multicast CAC.

## Viewing multicast CAC statistics

Statistics are accessible through CLI and SNMP.

### Viewing CAC statistics using CLI

The command `mcacStats` enables viewing multicast CAC statistics.

**Usage** `mcacStats - [ k | m | s | c ]`

| Command element | Description                |
|-----------------|----------------------------|
| k               | Display output in Kbps     |
| m               | Display output in Mbps     |
| s               | Display stats summary only |
| c               | Clear stats counters.      |

The default option is `-k`. The options `-k` and `-m` can be used with the `-s` option provided `-s` is the first option.

### Viewing CAC statistics using SNMP

The `mcastcacstats.mib` allows you to view the CAC statistics using the SNMP administrative utility.

```
admin> snmpwalk 110.110.110.9 public mcastCacStatsGroup
totalMcastCacBundleBw.0 = Gauge32: 800
totalStaticMcastCacBw.0 = Gauge32: 400
totalDynMcastCacBw.0 = Gauge32: 600
availDynMcastCacBw.0 = Gauge32: 400
usedStaticMcastCacBw.0 = Gauge32: 400
usedDynMcastCacBw.0 = Gauge32: 300
unConfGrpDrop.0 = Counter32: 4
mcastAddress.234.1.2.3 = IPAddress: 234.1.2.3
mcastAddress.234.1.2.4 = IPAddress: 234.1.2.4
mcastAddress.234.1.2.5 = IPAddress: 234.1.2.5
mcastAddress.234.1.2.6 = IPAddress: 234.1.2.6
priority.234.1.2.3 = 1
priority.234.1.2.4 = 0
priority.234.1.2.5 = 0
priority.234.1.2.6 = 1
joinsAccepted.234.1.2.3 = Gauge32: 1
joinsAccepted.234.1.2.4 = Gauge32: 1
joinsAccepted.234.1.2.5 = Gauge32: 0
joinsAccepted.234.1.2.6 = Gauge32: 1
joinsDenied.234.1.2.3 = Gauge32: 0
joinsDenied.234.1.2.4 = Gauge32: 0
joinsDenied.234.1.2.5 = Gauge32: 2
joinsDenied.234.1.2.6 = Gauge32: 0
```

The following table describes the MIB objects in the mcastcacstats.mib file:

MIB objects in the mcastcacstats.mib file.

| <b>MIB object</b>     | <b>Description</b>                                                                                                    |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------|
| totalMcastCacBundleBw | Total bandwidth assigned to this Multicast CAC bundle.                                                                |
| totalStaticMcastCacBw | Total bandwidth used by priority channels in this Multicast CAC bundle.                                               |
| totalDynMcastCacBw    | Total bandwidth configured for dynamic channels in this Multicast CAC bundle.                                         |
| availDynMcastCacBw    | Total bandwidth available for dynamic channels in this Multicast CAC bundle.                                          |
| usedStaticMcastCacBw  | Actual bandwidth used by priority channels in this Multicast CAC bundle.                                              |
| usedDynMcastCacBw     | Actual bandwidth used by dynamic channels in this Multicast CAC bundle.                                               |
| unConfGrpDrop         | The number of joins dropped for unconfigured groups.                                                                  |
| cacGroupTable         | This table contains the list of multicast groups configured for CAC control.                                          |
| cacGroupEntry         | Each entry in the table represents a unique multicast group configured for CAC control and its associated statistics. |
| mcastAddress          | Multicast IP address subjected to CAC control.                                                                        |
| priority              | Indicates whether this is a priority channel.                                                                         |
| joinsAccepted         | The number of joins served successfully for this multicast group.                                                     |
| joinsDenied           | The number of joins dropped for this multicast group due to CAC control.                                              |

---

# PIM-SM v2 Configuration

# 10

|                                                                |       |
|----------------------------------------------------------------|-------|
| PIM-SM features supported with this software version . . . . . | 10-2  |
| Overview of PIM-SM configuration . . . . .                     | 10-2  |
| Sample PIM-SM system configuration . . . . .                   | 10-10 |
| Administrative tools for PIM-SM routing . . . . .              | 10-11 |

Traditional multicast routing mechanisms, such as Distance Vector Multicast Routing Protocol (DVMRP) or Multicast OSPF (MOSPF) are intended for use in regions where groups are widely represented or bandwidth is universally plentiful. If these traditional schemes are used when multicast receivers and senders are distributed sparsely across a wide area, data packets and membership reports are sent over many links that do not lead to receivers or senders. For this reason, the traditional schemes can be inefficient for use in wide area networks.

Protocol Independent Multicast Sparse Mode (PIM-SM v2) is designed to operate efficiently across wide area networks, where groups are sparsely distributed. Each multicast group has a shared tree through which receivers learn about (“rendezvous with”) sources. The rendezvous point (RP) is the root of this per-group shared tree. PIM-SM uses RPs and explicit join/prune messages instead of the broadcast and prune mechanism used by PIM-Dense Mode or DVMRP.

This implementation of PIM-SM follows the current IETF drafts for PIM-SM v2 (draft-ietf-pim-sm-v2-new-07.txt, March 2003) and candidate bootstrap router (C-BSR) functionality (draft-ietf-pim-sm-bsr-03.txt, February 2003).



**Note** Stinger IP DSLAMs support full interoperability with RFC 2362-compliant multicast routers that have not yet implemented the PIM-SM v2 IETF draft recommendations. However, Stinger IP DSLAMs do not support interoperability with systems running PIM-SM v1. In addition, PIM-SM cannot be used with multiple MBONE interfaces in a Stinger IP DSLAM.

## ***PIM-SM features supported with this software version***

PIM-SM is currently supported on Gigabit Ethernet and trunk interfaces. Stinger IP DSLAMs cannot currently operate as PIM Multicast Border Router (PMBR). With the current software, Stinger systems support PIM-SM functionality as shown in Table 10-1:

*Table 10-1. Current level of support for PIM-SM functionality*

| <b>PIM-SM capability</b>                   | <b>Support in this software version</b>                                                                                                                                                                                                                                                                                                                                      |
|--------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PIM-SM general purpose states              | (*,G) state                                                                                                                                                                                                                                                                                                                                                                  |
| PIM join/prune messages                    | Join/prune messages for the (*,G) PIM-SM state                                                                                                                                                                                                                                                                                                                               |
| Bootstrap router (BSR)                     | <p>Stinger IP DSLAMs can be configured to act as C-BSR and take part in the BSR election process.</p> <p>If the Stinger becomes the elected BSR, it sends bootstrap messages (BSMs) to 224.0.0.13 on its PIM-enabled interfaces.</p> <p>If a PIM neighbor comes up and Stinger is the designated router on that interface, Stinger sends the recent BSM to the neighbor.</p> |
| Designated routers (DR) and hello messages | <p>On each PIM-enabled interface, the Stinger IP DSLAM can be elected DR on the basis of hello priority.</p> <p>Currently, Stinger IP DSLAMs do not support register messages and do not have the capability of processing IGMP messages on the LAN interface. For that reason, the Stinger must not be elected DR on the LAN interface if the LAN supports IGMP hosts.</p>  |
| RP-group mapping                           | <p>Static configuration of active group range and their respective RPs.</p> <p>Dynamic RP-group mapping using the BSR router mechanism.</p>                                                                                                                                                                                                                                  |
| Data packet forwarding                     | <p>Switch to SPT is not supported.</p> <p>Data packets are not forwarded from one PIM interface to another. Data packets flow from PIM interfaces to the users only.</p>                                                                                                                                                                                                     |

## ***Overview of PIM-SM configuration***

To enable the system to act as a multicast router in a PIM domain, you must complete the following steps:

- 1 Enable multicast forwarding and the PIM protocol. If appropriate, specify a C-BSR configuration for the Stinger IP DSLAM to participate in BSR elections and act as BSR if elected. See “Enabling multicast and PIM” on page 10-3.

- 2 Configure static mappings between multicast groups and PIM RPs. This is recommended as a failsafe configuration. See “Configuring static mappings between groups and rendezvous points” on page 10-5.
- 3 Configure an IP-INTERFACE or CONNECTION profile, to enable the system to operate as a PIM router on the Gigabit Ethernet interface, a trunk interface, or both. See “Configuring PIM on Gigabit Ethernet or trunk interfaces” on page 10-6.

For examples that show a system configuration that includes all three steps, see “Sample PIM-SM system configuration” on page 10-10.

## **Enabling multicast and PIM**

A PIM BSR is a dynamically elected router within a PIM domain that is responsible for constructing the set of RPs and originating BSMs. A C-BSR is a PIM router configured to participate in the BSR election and to act as BSR if elected. One BSR is elected per PIM domain on the basis of highest priority and address.

For details about the `pim bsr` and `pim rp` commands, see “Administrative tools for PIM-SM routing” on page 10-11.

## **Overview of settings in the IP-GLOBAL profile**

Following are the parameters, shown with default values, for enabling PIM-SM and configuring its BSR capabilities in a Stinger IP DSLAM:

```
[in IP-GLOBAL]
multicast-forwarding = no

[in IP-GLOBAL:pim-options]
enable = no
cbsr-enable = no
cbsr-ip-address = 0.0.0.0
cbsr-priority = 0
cbsr-interval = 60
```

| Parameter            | Setting                                                                                                                                                                                                                                                                                                                 |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| multicast-forwarding | Enables or disables multicast forwarding. This parameter must be set to <b>yes</b> for PIM-SM operations. When you change the value to <b>yes</b> and write the profile, the multicast subsystem reads the values in the IP-GLOBAL profile and initiates the forwarding function.                                       |
| enable               | Enables or disables the PIM routing protocol systemwide. This parameter and the <b>multicast-forwarding</b> parameter must be set to <b>yes</b> to enable PIM.                                                                                                                                                          |
| cbsr-enable          | Enables or disables the BSR router mechanism. When set to <b>yes</b> , the Stinger IP DSLAM acts as a candidate BSR and takes part in electing a BSR in the PIM domain. With the <b>yes</b> setting, you must specify a <b>cbsr-ip-address</b> value. This setting is not used when <b>enable</b> is set to <b>no</b> . |

| Parameter       | Setting                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| cbsr-ip-address | Local IP address the Stinger IP DSLAM uses to send BSMs when <code>cbsr-enable</code> is set to <code>yes</code> . This setting is not used when <code>cbsr-enable</code> is set to <code>no</code> .                                                                                                                                                                                        |
| cbsr-priority   | BSR priority for the Stinger, from 0 (the default) to 255. The priority is used in the election of BSR. The system is more likely to be elected BSR with a higher priority value. To enable the system to exchange BSMs without becoming BSR, leave the default zero setting, or set a low numeric value. This setting is not used when <code>cbsr-enable</code> is set to <code>no</code> . |
| cbsr-interval   | Number of seconds, from 5 to 900, between transmission of BSMs. The default is to send BSMs every 60 seconds. This setting is not used when <code>cbsr-enable</code> is set to <code>no</code> .                                                                                                                                                                                             |

### Example showing BSR election and dynamic group-RP mappings

The following commands configure the Stinger IP DSLAM to act as C-BSR using the IP address of the Gigabit Ethernet interface (1.1.1.101 in this example):

```
admin> read ip-global
admin> set multicast-forwarding = yes
admin> set pim-options enable = yes
admin> set pim-options cbsr-enable = yes
admin> set pim-options cbsr-ip-address = 1.1.1.101
admin> write -f
```

The following command displays the BSR status immediately after writing the IP-GLOBAL profile:

```
admin> pim bsr
Stinger BSR State : PENDING_BSR
Details of CURRENT BSR:
BSR IP Address      : 0.0.0.0
BSR Interface       : 0
BSR Priority         : 0
BSR holdtime        : 0
BSR Current Frag Tag : 0
BSR HASH masklen    : 0
```

After the bootstrap interval has elapsed and the system has received a BSM from another router in the domain, a repeat of the `pim bsr` command shows that the Stinger has become candidate BSR. For example:

```
admin> pim bsr
Stinger BSR State : CANDIDATE_BSR
Details of CURRENT BSR:
BSR IP Address      : 1.1.1.10
BSR Interface       : 1
```



```
BSR Priority          : 110
BSR holdtime         : 112
BSR Current Frag Tag : 0
BSR HASH masklen     : 30
```

The following commands modify the IP-GLOBAL profile to specify the highest BSR priority for the Stinger IP DSLAM:

```
admin> set cbsr-priority = 255
admin> write -f
```

Following an exchange of BSMs, the Stinger is elected BSR. For example:

```
admin> pim bsr
Stinger BSR State : ELECTED_BSR
Details of CURRENT BSR:
BSR IP Address      : 1.1.1.101
BSR Interface       : 1
BSR Priority         : 255
BSR holdtime        : 57
BSR Current Frag Tag : 717
BSR HASH masklen    : 30
```

The system is also receiving group-RP mappings dynamically from other routers. For example, the following command displays both static (if any) and dynamic group-RP mappings:

```
admin> pim rp
```

| Group        | RP-Address | RPF neighbor | Priority | holdtime |
|--------------|------------|--------------|----------|----------|
| 224.0.0.0/8  | 1.1.1.3    | -            | 40       | 75:73    |
| 234.0.0.0/8  | 1.1.1.3    | -            | 40       | 75:73    |
| 234.3.0.0/16 | 1.1.1.3    | -            | 40       | 75:73    |
| 234.4.0.0/16 | 1.1.1.3    | -            | 40       | 75:73    |
| 234.5.0.0/16 | 1.1.1.3    | -            | 40       | 75:73    |
| 234.6.0.0/16 | 1.1.1.3    | -            | 40       | 75:73    |

Static Entries:

|             |          |          |   |   |
|-------------|----------|----------|---|---|
| 225.0.0.0/8 | 1.1.1.10 | 1.1.1.10 | - | - |
| 235.0.0.0/8 | 1.1.1.10 | 1.1.1.10 | - | - |

## ***Configuring static mappings between groups and rendezvous points***

For PIM-SM to operate properly, all routers in the domain must share the same set of group-to-RP mappings. You can statically configure associations between groups and RPs by using the PIM-GROUP-RP-MAPPING profile. These static mappings provide a basic interoperability mechanism if the automatic methods of obtaining mappings should fail.

Each PIM-GROUP-RP-MAPPING profile specifies a mapping between an RP (specified as a reachable IP address) and a range of multicast groups (specified as a group and mask). The system uses these mappings to determine an RP for a given group.

Following are the parameters, shown with their default settings, for configuring a static group-to-RP mapping:

```
[in PIM-GROUP-RP-MAPPING/""]  
name* = ""  
rp-address = 0.0.0.0  
group-address = 0.0.0.0/0  
group-mask = 0.0.0.0
```

| Parameter     | Setting                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| name          | Text string, up to 31 characters long, that names the mapping between a multicast group and the IP address of an RP.                                                                                                                                                                                                                                                                                                                            |
| rp-address    | IP address of the RP. The address must be reachable throughout the domain.                                                                                                                                                                                                                                                                                                                                                                      |
| group-address | A multicast group address (a class D IP address). You can specify a full group address or a group range. If you specify a prefix (such as /8 in the value 226.0.0.0/8), the system automatically updates the group-mask parameter with the appropriate decimal value (such as 255.0.0.0).<br><br>The combined group address and group mask must be unique in the system. You cannot write duplicate mappings for the same group or group range. |
| group-mask    | A mask to be applied to the group-address value to obtain the group prefix mapped to the specified RP. For example, a value of 255.0.0.0 indicates a one-octet group prefix. If no mask is specified, the default mask of 255.255.255.255 is applied.                                                                                                                                                                                           |

The following commands create a static mapping for multicast group 231.1.1.1:

```
admin> new pim-group-rp-mapping 231  
admin> set rp-address = 1.1.1.3  
admin> set group-address = 231.1.1.1  
admin> write -f
```

The next command shows the static mapping in the group-RP set for the system:

```
admin> pim rp  
Group          RP-Address      RPF neighbor    Priority    holdtime  
Static Entries:  
231.1.1.1/32    1.1.1.3         -                -          -
```

## **Configuring PIM on Gigabit Ethernet or trunk interfaces**

The `pim-options` subprofiles in an IP-INTERFACE or CONNECTION profile specify settings to enable PIM-SM, and specify the handling of PIM hello and join/prune messages on the interface.

All PIM routers send hello messages periodically on each PIM-enabled interface, and record the hello information received from each PIM neighbor. Hello messages allow

a router to learn about the neighboring PIM routers on the interface, and the priority field in these messages is used in DR election on the LAN interface.

## PIM options in the IP-INTERFACE and CONNECTION profiles

Following are the parameters, shown with default values, for enabling PIM on the Gigabit Ethernet interface. The listing for trunk interface follows.

```
[in IP-INTERFACE/{ { shelf-1 first-control-module 2 } 0 }]
multicast-allowed = no

[in IP-INTERFACE/{ { shelf-1 first-control-module 2 } 0 }:pim-options]
enable = no
hello-interval = 30
hello-holdtime = 105
hello-priority-option = yes
hello-priority = 1
join-prune-interval = 60
join-prune-holdtime = 210
lan-delay-option = yes
lan-delay = 5000
override-interval = 2500
```



**Note** You cannot enable PIM on a virtual Ethernet interface.

Following are the parameters, shown with default values, for configuring PIM on a trunk interface:

```
[in CONNECTION/"" ]
encapsulation-protocol = atm-circuit

[in CONNECTION/"" :atm-options]
atm1483type = aal5-llc
vpi = 0
vci = 35
nailed-group = 1

[in CONNECTION/"" :ip-options]
ip-routing-enabled = yes
remote-address = 0.0.0.0/0
local-address = 0.0.0.0
multicast-allowed = no

[in CONNECTION/"" :ip-options:pim-options]
enable = no
hello-interval = 30
hello-holdtime = 105
hello-priority-option = yes
hello-priority = 1
join-prune-interval = 60
join-prune-holdtime = 210
lan-delay-option = yes
lan-delay = 5000
override-interval = 2500
```

| Parameter              | Setting                                                                                                                                                                                                                                                                                                                           |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| encapsulation-protocol | Set to atm for MPOA terminating connections.                                                                                                                                                                                                                                                                                      |
| atm1483type            | Stinger IP DSLAMs support the two encapsulation methods for carrying routed PDUs in the payload field of AAL type 5, which are defined in RFC 2684, <i>Multiprotocol Encapsulation over ATM Adaptation Layer 5</i> .<br><br>The aal5llc setting indicates LLC encapsulation. The aal5vc setting indicates AAL5 VC multiplexing.   |
| vpi                    | VPI for the MPOA terminating PVC. For details, see the <i>Stinger ATM Configuration Guide</i> .                                                                                                                                                                                                                                   |
| vci                    | VCI for the MPOA terminating PVC. For details, see the <i>Stinger ATM Configuration Guide</i> .                                                                                                                                                                                                                                   |
| nailed-group           | Nailed-group number of the trunk interface. For details, see the <i>Stinger ATM Configuration Guide</i> .                                                                                                                                                                                                                         |
| ip-routing-enabled     | IP routing must be enabled (as it is by default) for MPOA terminating connections.                                                                                                                                                                                                                                                |
| remote-address         | IP address of the remote device, which can include a subnet specification. If it does not include a subnet mask, the router software in the Stinger IP DSLAM assumes a default subnet mask that is based on address class.                                                                                                        |
| local-address          | IP address assigned to the local side of a numbered-interface connection. This is a requirement for PIM-enabled trunk interfaces.                                                                                                                                                                                                 |
| multicast-allowed      | Must be set to yes for PIM-enabled interfaces.                                                                                                                                                                                                                                                                                    |
| enable                 | Enables or disables the PIM routing protocol on the interface.                                                                                                                                                                                                                                                                    |
| hello-interval         | Number of seconds between sending hello messages to PIM neighbors on this interface. The valid range is from 1 to 65535 with a default value of 30 seconds. The value must be less than that of the hello-holdtime parameter.                                                                                                     |
| hello-holdtime         | Number of seconds a receiver of hello messages must consider the sender reachable before timing out the sender. The valid range is from 1 to 65535 with a default value of 105 seconds. The value must be greater than that of the hello-interval parameter.                                                                      |
| hello-priority-option  | Whether the Stinger IP DSLAM will participate in DR election on this interface (yes or no, with a default value of yes).                                                                                                                                                                                                          |
| hello-priority         | DR election priority for the Stinger IP DSLAM on the trunk interface. The DR election priority is a 32-bit unsigned number contained in a hello message. A router with a numerically larger priority is preferred in electing a new DR. The valid range for this setting is from 0 to 4,294,967,295, with a default setting of 1. |

| Parameter           | Setting                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| join-prune-interval | <p>Number of seconds between sending PIM join/prune messages to PIM neighbors on this interface. A join/prune message consists of a list of groups and a list of joined and pruned sources for each group. The valid range is from 1 to 65535 with a default value of 60 seconds. The value must be less than that of the join-prune-holdtime parameter.</p> <p><b>Note</b> Stinger IP DSLAMs do not currently support the (S,G) state, so it always sends (*,G) join/prune messages.</p> |
| join-prune-holdtime | <p>Number of seconds a receiver of join/prune messages must consider the list valid before timing out the information. The valid range is from 1 to 65535 with a default value of 210 seconds. The value must be greater than that of the join-prune-interval parameter.</p>                                                                                                                                                                                                              |
| lan-delay-option    | <p>Whether the Stinger IP DSLAM will expect propagation delay over an Ethernet interface (yes or no, with a default value of yes). The lan-delay option is not sent in hello messages on a trunk interface, even if it is configured, because it applies only to LAN interfaces.</p>                                                                                                                                                                                                      |
| lan-delay           | <p>Number of milliseconds of expected propagation delay over the Ethernet interface. The valid range is from 1 to 65535 with a default value of 5000 milliseconds. This value is not sent in hello messages on a trunk interface, even if it is configured, because it applies only to LAN interfaces.</p>                                                                                                                                                                                |
| override-interval   | <p>A delay interval, in milliseconds, used to randomize when scheduling a delayed join message. The valid range is from 1 to 65535 with a default value of 2500 milliseconds. Because the override-interval is sent along with lan-delay, this value is not sent in hello messages on a trunk interface, even if it is configured.</p>                                                                                                                                                    |

### Example of enabling PIM on the Gigabit Ethernet interface

The following commands enable PIM on a Gigabit Ethernet interface:

```
admin> read ip-interface { { 1 8 2 } 0 }  
admin> set ip-address = 1.1.1.2  
admin> set multicast-allowed = yes  
admin> set pim-options enable = yes  
admin> write -f
```

## Example of enabling PIM on a trunk interface


The following commands enable PIM on an MPOA terminating PVC on a trunk port. A numbered interface is required. For details about numbered interfaces, see “Example of using a local-address setting for a numbered interface” on page 4-52.

```
admin> new connection pim-trunk
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 123.123.123.2/24
admin> set ip-options local-address = 123.123.123.1/24
admin> set ip-options multicast-allowed = yes
admin> set ip-options pim-options enable = yes
admin> set telco-options nailed-groups = 851
admin> set mp-options enabled = no
admin> set atm-options vci = 38
admin> set atm-options nailed-group = 851
admin> write -f
```

## Sample PIM-SM system configuration

For details about multicast client configuration, see “Configuring multicast client interfaces” on page 9-11. In the sample setup shown in Figure 10-1, PIM-SM is enabled on the Gigabit Ethernet port.

Figure 10-1. PIM-SM on Gigabit Ethernet and trunk interface



The following commands enable PIM-SM and configure the Stinger IP DSLAM to act as C-BSR using the IP address of the Gigabit Ethernet interface (1.1.1.101 in this example):

```
admin> read ip-global
admin> set multicast-forwarding = yes
admin> set pim-options enable = yes
admin> set pim-options cbsr-enable = yes
admin> set pim-options cbsr-ip-address = 1.1.1.101
admin> write -f
```

The following commands configure a PVC for multicast client CPE router in Figure 10-1:

```
admin> new connection mcast-client
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 2.2.2.2/32
admin> set ip-options multicast-allowed = yes
admin> set ip-options multicast-rate-limit = 20
admin> set atm-options vci = 100
admin> set atm-options nailed-group = 51
admin> write -f
```

The following commands enable multicast and PIM-SM on a Gigabit Ethernet interface:

```
admin> read ip-interface { { 1 8 2 } 0 }
admin> set ip-address = 1.1.1.101
admin> set multicast-allowed = yes
admin> set pim-options enable = yes
admin> write -f
```

The following commands configure a group-to-RP mapping specifying the RP shown in Figure 10-1:

```
admin> new pim-group-rp-mapping 229
admin> set rp-address = 1.1.1.10
admin> set group-address = 229.0.0.0/8
admin> write -f
```

Following is the relevant group mapping:

```
admin> pim rp
Group          RP-Address      RPF neighbor    Priority    holdtime
Static Entries:
229.0.0.0/8    1.1.1.10        -                -           -
```

The following command shows the PIM neighbor across the Gigabit Ethernet interface:

```
admin> pim nbr
Neighbor      Interface    Priority    Holdtime      DR
1.1.1.10      1            100        105:96        No
```

## ***Administrative tools for PIM-SM routing***

The `pim` command displays PIM-related information for active PIM-enabled interfaces in the system. For details, see “PIM-SM diagnostics” on page A-10.

The system provides SNMP MIB support for the PIM protocol as defined in `draft-ietf-pim-mib-v2-01.txt` (the PIMv2 MIB). The PIMv2 MIB is placed in the MIB tree under `experimental 61`. For details, see “PIMv2 MIB support” on page A-35.

In addition, the output of the `netstat -s` command now includes the total PIM statistics for all PIM-enabled interfaces in the system. For example:

```
admin> netstat -s
```

```
...
```

```
pim:
```

```
    25 packets received
```

```
    24 hello packet received
```

```
    1 C-RP packets received
```

```
    38 packets transmitted
```

```
    26 hello packets sent
```

```
    12 Bootstrap packets sent
```



---

# Filter Configuration

# 11

|                                             |       |
|---------------------------------------------|-------|
| Filter overview . . . . .                   | 11-1  |
| Defining IP filters . . . . .               | 11-2  |
| Defining ICMP filters . . . . .             | 11-12 |
| Defining route filters . . . . .            | 11-17 |
| Defining Ethernet input filters . . . . .   | 11-19 |
| Applying a filter to an interface . . . . . | 11-21 |
| Administrative tools for filters . . . . .  | 11-23 |

In a FILTER profile, a Stinger IP DSLAM currently supports `ip-filter`, `route-filter`, and `ethernet-filter` subprofile configurations. ICMP packet filters are configured within the `ip-filter` subprofile. In addition, the `gen-filter` subprofile has one supported use—to define an explicit default rule for an IP filter.

## Filter overview

A Stinger IP DSLAM supports the following types of filter:

- IP filters, for filtering packets based on IP packet header fields.
- Route filters, for filtering specific routes in RIP update packets.
- Ethernet input filters, for filtering inbound packets on the basis of Ethernet protocol or MAC addresses (or both).

After you define a filter and apply it to an interface, the system monitors traffic on the interfaces and, when a match occurs, takes the specified action.

Table 11-1. Default filtering behavior

| Filters                                     | Action on packets                                                                                                                                                                                   |
|---------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| No filter is applied to the connection.     | No packets are dropped due to filtering.                                                                                                                                                            |
| A filter rule is applied to the connection. | All packets except those matching the filter rule are dropped. Packets that match the filter rule are forwarded or dropped according to the definition in the rule (forward = yes or forward = no). |

Table 11-1. Default filtering behavior (Continued)

| Filters                                                | Action on packets                                                                                                                                                                         |
|--------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| An explicit default rule is applied to the connection. | All packets that do not match specific filter rules are forwarded or dropped according to the definition in the default rule ( <code>forward = yes</code> or <code>forward = no</code> ). |

## Filter rules

A FILTER profile includes up to 30 input filters and 30 output filters, each of which specifies a set of rules in one direction (inbound or outbound). Rules define comparisons that are made in a defined order and an action to be taken when a match occurs.

The system applies the rules in sequence, from 1 to 30. The action specified in the rule can apply to packets that match the rule, or to all packets *except* those that match the rule.

The filtering process stops immediately if conditions in a packet match a rule, and the action in that rule is applied to the packet. If a packet does not match a rule, the next rule is applied to the packet, and so forth. If no matches occur, the filtering subsystem drops the nonmatching packet. However, if an explicit default rule is specified as the last rule in the sequence, the system applies the action specified in that rule.

## Explicit default filter rules

For filtered connections, the system default is to drop all packets that do not match the filter rules. If an explicit default rule is defined, the action specified in that rule overrides the system default. An explicit default rule has following settings:

- All filter fields of the default rule must be set to 0 (default values when the filter is created).
- The `valid-entry` field must be set to `yes`
- The `forward` entry must be set to the desired default action (`drop` or `forward`).
- The `type` of the filter must be set.

If a filter specifies more than one explicit default rule, only the first one is taken into account.

## Defining IP filters

IP filters use IP packet header fields to select traffic to be filtered. You apply an IP filter to an IP interface.

## Overview of IP filter settings

Following are the FILTER profile settings for defining IP filters. The parameters are shown with their default values for input filter rules. The same values apply for output filter rules—setting the parameters in an input filter affects the inbound data stream, and setting them in an output filter rule affects the outbound data stream.

```
[in FILTER/""]  
filter-name* = ""  
  
[in FILTER/":input-filters[1]]  
valid-entry = no  
forward = no  
type = gen-filter  
  
[in FILTER/":input-filters[1]:ip-filter]  
protocol = 0  
source-address-mask = 0.0.0.0  
source-address = 0.0.0.0  
dest-address-mask = 0.0.0.0  
dest-address = 0.0.0.0  
Src-Port-Cmp = none  
source-port = 0  
Dst-Port-Cmp = none  
dest-port = 0  
tcp-estab = no  
icmp-type = any
```

| Parameter           | Setting                                                                                                                                                                                                          |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| filter-name         | Name of the FILTER profile, up to 36 characters. You apply a filter to an interface by referring to this name.                                                                                                   |
| valid-entry         | Enable/disable the input or output rule. With a setting of <b>no</b> (the default), the system skips this rule when filtering the data stream. Set this parameter to <b>yes</b> for each rule you define.        |
| forward             | Forwarding action for the rule. The default value of <b>no</b> causes the system to discard matching packets.                                                                                                    |
| type                | Type of filter ( <b>ip-filter</b> , <b>gen-filter</b> , <b>route-filter</b> , and <b>ethernet-filter</b> ). Only the parameters in the corresponding subprofile are applicable for the rule.                     |
| protocol            | Protocol number. A number of 0 (zero) matches all protocols. A nonzero number is compared to the Protocol field in each packet. For a list of assigned protocol numbers, see RFC 1700, <i>Assigned Numbers</i> . |
| source-address-mask | Mask to be applied to the <b>source-address</b> value before comparing that value to the source address of a packet.                                                                                             |
| source-address      | IP address. After applying the <b>source-address-mask</b> value, the system compares the result to the source address in a packet. See “Filtering on source or destination IP addresses” on page 11-5.           |
| dest-address-mask   | A mask to be applied to the <b>dest-address</b> value before comparing that value to the destination address of a packet.                                                                                        |
| dest-address        | IP address. After applying the <b>dest-address-mask</b> value, the system compares the result to the destination address in a packet. See “Filtering on source or destination IP addresses” on page 11-5.        |

| Parameter    | Setting                                                                                                                                                                                                                                                                     |
|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| src-port-cmp | Type of comparison to be done on the source UDP/TCP port. The <b>less</b> (less than) and <b>gtr</b> (greater than) operators are not supported when comparing source port values in traffic destined for an external system. See “Filtering on port numbers” on page 11-5. |
| source-port  | Port number to be compared against the source port of a packet.                                                                                                                                                                                                             |
| dst-port-cmp | Type of comparison to be done on the destination UDP/TCP port. See “Filtering on port numbers” on page 11-5.                                                                                                                                                                |
| dest-port    | Port number to be compared against the destination port of a packet.                                                                                                                                                                                                        |
| tcp-estab    | TCP established session. For details, see “TCP-established filters” on page 11-6.                                                                                                                                                                                           |
| icmp-type    | Type of ICMP message to filter. See “Defining ICMP filters” on page 11-12.                                                                                                                                                                                                  |

### ***Details of IP filter comparison passes***

The system compares a packet to rule #1. If the comparison fails (the packet does not match the rule), the system proceeds to rule #2, and so forth. When a comparison succeeds, the filtering process stops and the forwarding action in that filter is applied to the packet. Within each input or output filter, the comparisons proceed as follows:

- 1 Apply the **source-address-mask** value to the **source-address** value and compare the result to the source address of the packet. If they are not equal, the comparison fails.
- 2 Apply the **dest-address-mask** value to the **dest-address** value and compare the result to the destination address in the packet. If they are not equal, the comparison fails.
- 3 If the **protocol** parameter is zero (which matches any protocol), the comparison succeeds. If it is nonzero and not equal to the protocol field in the packet, the comparison fails. If port numbers are being used in the rule, either alone or with IP addresses, the **protocol** parameter must be set to the corresponding value.
- 4 If the **src-port-cmp** parameter is not set to **none**, compare the **source-port** number to the source port number of the packet. If they do not match as specified by the **src-port-cmp** parameter, the comparison fails.
- 5 If the **dst-port-cmp** parameter is not set to **none**, compare the **dest-port** number to the destination port number of the packet. If they do not match as specified by the **dst-port-cmp** parameter, the comparison fails.

If all comparisons fail, the packet does not match the filter. For security purposes, the Stinger IP DSLAM does not automatically forward nonmatching packets unless the filter explicitly allows nonmatching packets to pass.

## Filtering on source or destination IP addresses

When you specify a source or destination address in an IP filter, the system applies the filter's forwarding action to packets received from or sent to that address. If you also specify a subnet mask, the system applies the mask to the address value before comparing the resulting value to the source or destination address in a packet.

To apply the mask, the system translates both the mask and address values into binary format and then uses a logical AND to apply the mask to the address. The mask hides the bits whose positions match those of the binary zeros in the mask. A mask of all zeros (the default) masks all bits. If the address value itself is also all zeros (the default), the filter matches any source or destination address. A mask of all ones (255.255.255.255) masks no bits, so the full address for a single host is compared to the address value.

You can use the address mask to mask out the host portion of an address, for example, or the host and subnet portion, so the rule matches the address to or from any host on a given network.

## Filtering on port numbers

IP filters can specify a port number to be compared to the source or destination port (or both) in a packet. A port number of zero matches nothing. TCP and UDP port numbers are typically assigned to services. For a list of well-known port assignments, see RFC 1700, *Assigned Numbers*.



**Note** For security purposes, Alcatel-Lucent recommends that you filter all services from outside your domain that are not required. UDP-based services make your network particularly vulnerable to certain types of security attacks.

The type of comparison determines when a match occurs. For source port values, filters applied to traffic destined for an external system support the **none** (no comparison is made) or **eq** (equal to) operators.

For other traffic, the following operators can be used to compare source port or destination port values:

- **none** (no comparison is made)
- **eq** (equal to)
- **less** (less than)
- **gtr** (greater than)



**Note** The **neq** (not equal to) operator is not supported for port comparisons.

The following commands show an *illegal* rule that uses the unsupported **neq** operator to forward packets with a source port not equal to 50:

```
admin> set input-filters 1 valid-entry = yes
admin> set input-filters 1 Type = ip-filter
admin> set input-filters 1 forward = yes
admin> set input-filters 1 ip-filter protocol = 17
admin> set input-filters 1 ip-filter Src-Port-Cmp = neq
admin> set input-filters 1 ip-filter source-port = 50
```

If the `src-port-cmp` rule is set to the unsupported value (`neq`), the system logs the following warning message:

```
LOG warning, Shelf 1, Controller-1, Time: 02:23:31--
```

```
IP Filters: Not equal operation not supported for source port comparison
```

In this case, the faulty rule specifying the `neq` operator is not applied. The other rules of the filter are applied to the traffic stream.

The following commands show a legal workaround using the `less` and `gtr` comparison operator in two rules to accomplish the same effect as using the unsupported `neq` operator:

```
admin> set input-filters 1 valid-entry = yes
admin> set input-filters 1 Type = ip-filter
admin> set input-filters 1 forward = yes
admin> set input-filters 1 ip-filter protocol = 17
admin> set input-filters 1 ip-filter Src-Port-Cmp = less
admin> set input-filters 1 ip-filter source-port = 50
admin> set input-filters 2 valid-entry = yes
admin> set input-filters 2 Type = ip-filter
admin> set input-filters 2 forward = yes
admin> set input-filters 2 ip-filter protocol = 17
admin> set input-filters 2 ip-filter Src-Port-Cmp = gtr
admin> set input-filters 2 ip-filter source-port = 50
```

## TCP-established filters

You can define output filters that allow TCP traffic only on established connections and deny all TCP traffic that is not part of an established session, such as traffic sourced from the external network.

This type of filter is supported only for `output-filters`. It is not supported for `input-filters`. Further, it is recommended that you configure this type of filter in the *first* indexed `output-filters` subprofile (`output-filters[1]`), because the `tcp-estab` field is always the first filter rule applied, even if it is configured using a higher index.

This type of IP filter affects IP routed traffic only. Note also that filters are not applied to L2TP-tunneled traffic.

### *How TCP-established filters work*

Typically, an IP filter contains several specifications that are compared to packet header fields, such as `source-ip-address`, `destination-ip-address`, `protocol`, and so forth. To have a match with the filter, *all* specified fields must match.

The TCP-established filters, the `tcp-estab` field and the `protocol` field are the only relevant specifications in the `ip-filter` subprofile. All other settings in the `ip-filter` subprofile *must* have their default values.

As a result, *all* TCP packets are filtered solely on the basis of the Acknowledgment (ACK) or Reset (RST) bit (or both) in the TCP header. So, applying a TCP-established filter to a WAN connection causes the system to discard all TCP session initiation from the Internet to the DSL client.

## Overview of required settings

Following are the relevant parameters, shown with default settings. These are the only parameters to be configured for this type of filter.

```
[in FILTER/"":output-filters[1]]
valid-entry = no
forward = no
Type = gen-filter

[in FILTER/"":output-filters[1]:ip-filter]
protocol = 0
tcp-estab = no
```

| Parameter           | Setting                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| valid-entry         | Enables or disables the output filter. Set this parameter to <b>yes</b> in <code>output-filters[1]</code> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| forward             | Forwarding action for the rule. Keep the default value of <b>no</b> to drop packets that are not initiated in an established session.                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| type                | Type of filter. Set this parameter to <b>ip-filter</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| ip-filter:protocol  | Protocol number. Set this parameter to 6 for TCP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| ip-filter:tcp-estab | TCP established session. When set to <b>yes</b> , the system examines the TCP-options field of a packet's TCP header to verify that the Acknowledgment (ACK) or Reset (RST) bit is set. These fields indicate whether the packets belong to an existing connection.<br><br>If the rest of the filter definition is configured as shown in "Sample TCP-established filter" on page 11-8, the system drops all traffic that is not initiated from an established connection. Applying this filter as an output filter of WAN connections will discard all TCP session initiation from the Internet to the DSL client. |

If a filter has the `tcp-estab` parameter set to **yes**, but one or more of the other parameters in the filter does not have the required value, the output filter is skipped and a **WARNING** message is logged stating that the `tcp-estab` parameter was not applied due to a faulty configuration. For example:

```
8/01/2005 10:48:38 shelf-1/first-contro warning tcp-estab parameter
not applied for i/f 1 2 15, (Filter not properly configured).
```

## Sample IP filters

The following sections present sample IP filters and illustrate some of the issues you might consider when writing your own IP filters. The sample filters presented here do not address the fine points of network security. You might want to use these filters as a starting point and augment them to address your security requirements.



**Note** Some of the sample IP filters in the next sections include explicit default rules of type `ip-filter`. In general, an explicit default rule of type `generic-filter` is

recommended to ensure the expected behavior in all types of profiles. For an example, see “Sample filter using a generic explicit default rule” on page 11-11.

### Sample TCP-established filter

The following commands configure a filter that drops all TCP packets that are not part of an established connection:

```
admin> new filter estab-only
admin> set output-filters 1 valid-entry = yes
admin> set output-filters 1 type = ip-filter
admin> set output-filters 1 ip-filter protocol = 6
admin> set output-filters 1 ip-filter tcp-estab = yes
admin> set output-filters 2 valid-entry = yes
admin> set output-filters 2 forward = yes
admin> write -f
```

The next commands configure a BIR connection that applies the filter:

```
admin> new connection bir-user
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 192.168.150.2/24
admin> set ip-options local-address = 192.168.150.1/24
admin> set session-options data-filter = estab-only
admin> set atm-options nailed-group = 51
admin> write -f
```

When the connection is actively passing data, you can check for packet discards by using the `stats error np fpp` command. For example, the following command output indicates that 186 output packets have been discarded:

```
admin> stats error np fpp
Pass#2 Error:
-----
IN Filter Discard           :          0
OUT Filter Discard          :         186
```

### Preventing address spoofing

Spoofing IP packets allows an intruder on a remote network to impersonate a local system's IP address. This section presents an example of an IP filter that prevents address spoofing when applied to a CPE interface.

The following commands create input filter #, which drops packets that have a local source address. In this example, the local network has an IP address of 192.100.50.128, with a subnet mask of 255.255.255.192. These values are just arbitrary examples. Because `forward` is set to `no` (the default), inbound packets with a source address on the LAN will be dropped.

```
admin> new filter ip-spoof
admin> set input 1 valid = yes
```



```
admin> set input 1 type = ip-filter
admin> set input 1 ip-filter source-address-mask = 255.255.255.192
admin> set input 1 ip-filter source-address = 192.100.50.128
```

The next set of commands creates input filter #2, which drops packets with a source address equal to the loopback address (127.0.0.0).

```
admin> set input 2 valid = yes
admin> set input 2 type = ip-filter
admin> set input 2 ip-filter source-address-mask = 255.0.0.0
admin> set input 2 ip-filter source-address = 127.0.0.0
```

The next set of commands creates input filter #3, which explicitly accepts all remaining source addresses and forwards them to the local network. Except for `forward = yes`, the third filter uses all default values. Because `forward` is set to `yes`, the system forwards all remaining packets (those with nonlocal source addresses) to the LAN.

```
admin> set input 3 valid = yes
admin> set input 3 forward = yes
admin> set input 3 type = ip-filter
```

The next set of commands creates an output filter and sets the forwarding action to `yes`. This filter specifies the source mask and address for the local network. (Packets originating on the local network should be forwarded to the CPE.)

```
admin> set output 1 valid = yes
admin> set output 1 type = ip-filter
admin> set output 1 forward = yes
admin> set output 1 ip-filter source-address-mask = 255.255.255.192
admin> set output 1 ip-filter source-address = 192.100.50.128
admin> write -f
```

## An IP filter for more complex security issues

In this example, the local network supports a Web server, and the administrator needs to carry out the following tasks:

- Provide client access to the server's IP address.
- Restrict ingress traffic to all other hosts on the local network.

However, many local IP hosts need to access the Internet and use IP-based applications such as `telnet` or `ftp`, so their response packets need to be directed appropriately to the originating host. In this example, the Web server's IP address is 192.9.250.5. The filter will be applied in CONNECTION profiles as a data filter.

The following commands create the first input filter, which sets `forward` to `yes` and allows packets to reach the Web server's destination address at a destination TCP port that can be used for `telnet` or `ftp`:

```
admin> new filter web-access
admin> set input 1 valid = yes
admin> set input 1 forward = yes
```

```
admin> set input 1 type = ip-filter
admin> set input 1 ip-filter protocol = 6
admin> set input 1 ip-filter dest-address-mask = 255.255.255.255
admin> set input 1 ip-filter dest-address = 192.9.250.5
admin> set input 1 ip-filter dst-port-cmp = eql
admin> set input 1 ip-filter dest-port = 80
```

The next set of commands creates the second input filter, which allows inbound TCP packets in response to a local user's outbound telnet request, by specifying that TCP packets whose destination port number is higher than that of the source port are forwarded. (The telnet requests go out on port 23, and responses come back on some random port above port 1023.)

```
admin> set input 2 valid = yes
admin> set input 2 forward = yes
admin> set input 2 type = ip-filter
admin> set input 2 ip-filter protocol = 6
admin> set input 2 ip-filter dst-port-cmp = gtr
admin> set input 2 ip-filter dest-port = 1023
```

The next set of commands creates the third input filter, which allows inbound RIP updates, by specifying that inbound UDP packets are forwarded if the destination port number is higher than that of the source port. (For example, suppose a RIP packet goes out as a UDP packet to destination port 520. The response to this request goes to a random destination port above port 1023.)

```
admin> set input 3 valid = yes
admin> set input 3 forward = yes
admin> set input 3 type = ip-filter
admin> set input 3 ip-filter protocol = 17
admin> set input 3 ip-filter dst-port-cmp = gtr
admin> set input 3 ip-filter dest-port = 1023
```

The following commands create the fourth input filter. This filter uses all default values, which allows unrestricted use of ping and traceroute.

```
admin> set input 4 valid = yes
admin> set input 4 forward = yes
admin> set input 4 type = ip-filter
admin> write -f
```

### Sample filter with no explicit default rule

When you do not define an explicit default rule, the system uses the implicit default which discards all packets that do not match the input and output filter rules. If a filter defines rules only in one direction, traffic in the other direction is not filtered. For example, the following sample filter specifies that UDP packets from a source port less than 50 should be forwarded. Because it does not specify an explicit default rule, this filter causes all packets that do not match input filter 1 to be dropped.

```
admin> new filter input-filter-1
admin> set input-filters 1 valid-entry = yes
admin> set input-filters 1 forward = yes
admin> set input-filters 1 Type = ip-filter
admin> set input-filters 1 ip-filter protocol = 17
admin> set input-filters 1 ip-filter source-address-mask = 255.255.255.255
admin> set input-filters 1 ip-filter source-address = 192.168.2.2
admin> set input-filters 1 ip-filter Src-Port-Cmp = less
admin> set input-filters 1 ip-filter source-port = 50
admin> write -f
```

### Sample filter with explicit default rule

Defining an explicit default rule allows more flexibility in specifying which packets to drop. For example, the following sample filter specifies that UDP packets from a source port less than 50 should be *dropped*, and includes an explicit default rule that causes all other incoming IP packets to be *forwarded*. Because this filter does not define output rules, all packets are forwarded in the output direction.

```
admin> new filter input-filter-2
admin> set filter-name = input-filter
admin> set input-filters 1 valid-entry = yes
admin> set input-filters 1 forward = no
admin> set input-filters 1 Type = ip-filter
admin> set input-filters 1 ip-filter protocol = 17
admin> set input-filters 1 ip-filter source-address-mask = 255.255.255.255
admin> set input-filters 1 ip-filter source-address = 192.168.2.2
admin> set input-filters 1 ip-filter Src-Port-Cmp = less
admin> set input-filters 1 ip-filter source-port = 50
admin> set input-filters 2 valid-entry = yes
admin> set input-filters 2 forward = yes
admin> set input-filters 2 Type = ip-filter
admin> write -f
```

### Sample filter using a generic explicit default rule



**Note** A Stinger IP DSLAM does not support generic packet filters, but you can specify a generic explicit default rule in an IP filter. This is the only supported use of **generic-filter** for connections terminating in the Stinger IP DSLAM. A generic default rule affects *all* packet types that do not match the filter rules for a certain direction (input or output).

To define an explicit default rule that affects *all* packets that do not match the filter rules in a certain direction, set the **Type** value to **generic-filter**.

If the explicit default filter rule is of type **ip-filter**, ARP packets (and other non-IP packets) will not be affected by the default. For example, if you want to forward all

nonmatching packets including non-IP packets such as ARP packets, you must create a generic-filter rule as default with the action set to forward.

The following filter specifies a generic explicit default rule to allow forwarding of all incoming packets that do not match the input filter rules. Because the output direction does not specify an explicit default rule, all packets that do not match the output filter rules will be dropped.

```
admin> new filter input-output
admin> set input-filters 1 valid-entry = yes
admin> set input-filters 1 forward = no
admin> set input-filters 1 Type = ip-filter
admin> set input-filters 1 ip-filter protocol = 17
admin> set input-filters 1 ip-filter source-address-mask = 255.255.255.255
admin> set input-filters 1 ip-filter source-address = 192.168.2.2
admin> set input-filters 1 ip-filter Src-Port-Cmp = less
admin> set input-filters 1 ip-filter source-port = 50
admin> set input-filters 2 valid-entry = yes
admin> set input-filters 2 forward = yes
admin> set input-filters 2 Type = generic-filter
admin> set output-filters 1 valid-entry = yes
admin> set output-filters 1 forward = yes
admin> set output-filters 1 Type = ip-filter
admin> set output-filters 1 ip-filter protocol = 17
admin> set output-filters 1 ip-filter Src-Port-Cmp = less
admin> set output-filters 1 ip-filter source-port = 50
admin> write -f
```

## **Defining ICMP filters**

Filtering ICMP packets by message type enables network administrators to send OAM traffic downstream and receive the required response, while preventing DSL users from initiating OAM traffic to the DSLAM or upstream network element.

You can define ICMP packet filtering in a general way, for example, to prevent DSL users from issuing any ICMP Echo-Request packets on their connections. Or you can create filters that prevent users from pinging specific IP addresses, or to prevent Echo-Requests from specific addresses from being forwarded upstream.

## **Overview of profile settings**

To filter on ICMP message type, you must set the `protocol` and `icmp-type` parameters in the `ip-filter` subprofile. You can optionally also set the source or destination address (or both).

The parameters are shown here with default values for an output filter. No other parameters in the `ip-filter` subprofiles apply to ICMP filtering.

```
[in FILTER/"" :output-filters[1]:ip-filter]
protocol = 0
source-address-mask = 0.0.0.0
source-address = 0.0.0.0
dest-address-mask = 0.0.0.0
dest-address = 0.0.0.0
icmp-type = any
```

| Parameter            | Setting                                                                 |
|----------------------|-------------------------------------------------------------------------|
| protocol             | ICMP filtering applies only when protocol is set to 1.                  |
| source-address       | Source address to compare.                                              |
| source-address-mask  | Netmask for source address.                                             |
| dest-address         | Destination address to compare.                                         |
| dest-address-mask    | Netmask for destination address.                                        |
| icmp-type            | The ICMP message type to filter.                                        |
| any                  | Matches any ICMP message type.<br>This is the default setting.          |
| echo-request         | RFC 792 Type 8, Echo request.                                           |
| echo-reply           | RFC 792 Type 0, Echo reply.                                             |
| dest-unreachable     | RFC 792 Type 3, Destination unreachable.                                |
| src-quench           | RFC 792 Type 4, Source quench.                                          |
| redirect             | RFC 792 Type 5, Redirect.                                               |
| router-advertisement | RFC 1256 Type 9, Router advertisement.                                  |
| router-solicitation  | RFC 1256 Type 10, Router solicitation. ('Router selection' in RFC 1256) |
| time-exceeded        | RFC 792 Type 11, Time exceeded.                                         |
| param-problem        | RFC 792 Type 12, Parameter problem.                                     |
| time-request         | RFC 792 Type 13, Timestamp request.                                     |
| time-reply           | RFC 792 Type 14, Timestamp reply.                                       |
| info-request         | RFC 792 Type 15, Information request.                                   |
| info-reply           | RFC 792 Type 16, Information reply.                                     |

| Parameter       | Setting                                                                  |
|-----------------|--------------------------------------------------------------------------|
| netmask-request | RFC 950 Type 17, Netmask request.<br>(‘Address Mask Request’ in RFC 950) |
| netmask-reply   | RFC 950 Type 18, Netmask reply.<br>(‘Address Mask Reply’ in RFC 950)     |

## Sample ICMP filter configurations

This section contains the following sample configurations:

- A filter for blocking incoming Echo-Request packets on one or more DSL user connections
- A filter for blocking Echo-Request packets from being forwarded to an upstream network element, such as a BRAS
- A filter for blocking Echo-Request packets to a specific address

### Blocking incoming Echo-Request packets

In the following example, a filter is defined to drop inbound ICMP Echo-Request packets, and the filter is applied to DSL user connections. Upstream network users are still able to send ICMP Echo-Request packets to the DSL users, and to receive the ICMP Echo-Replies.

- 1 Define the filter to drop inbound ICMP Echo-Request packets.  

```
admin> new filter in-no-echo-req
admin> set input-filters 1 valid-entry = yes
admin> set input-filters 1 Type = ip-filter
admin> set input-filters 1 ip-filter protocol = 1
admin> set input-filters 1 ip-filter icmp-type = echo-request
admin> set input-filters 2 valid-entry = yes
admin> set input-filters 2 forward = yes
admin> write -f
```
- 2 Configure an IPoA connection for a DSL user and apply the filter.  

```
admin> new connection ipoa-11-40
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 10.40.10.1/24
admin> set session-options data-filter = in-no-echo-req
admin> set atm-options nailed-group = 540
admin> write -f
```
- 3 Configure a BIR connection for a DSL user and apply the same filter.  

```
admin> new connection bir-11-37
admin> set active = yes
admin> set encapsulation-protocol = atm
```

```
admin> set ip-options remote-address = 10.10.10.0/24
admin> set ip-options local-address = 10.10.10.3/32
admin> set session-options data-filter = in-no-echo-req
admin> set atm-options nailed-group = 537
admin> set bir-options enable = yes
admin> set bir-options proxy-arp = yes
admin> write -f
```

Incoming Echo-Request packets on these connections will now be dropped. All other traffic, including Echo-Replies and all data traffic, will be forwarded.

You can use the `filterdisp` command to display details about the filter application. For example:

```
admin> filterdisp
ID   Username   Src   Route-Filter  Data-Filter  Call-Filter  TOS-Filter
-----
015  bir-11-37   loc                in-no-echo-re
019  ipoa-11-40  loc                in-no-echo-re
<end user list> 2 active user(s)
```

```
admin> filterdisp 19
Hostname:      ipoa-11-40
```

```
*****
```

```
Data Filter      Direction: In
```

```
-----
```

```
Forward = no
Type = IP Filter
protocol = 1
source-address-mask = 0.0.0.0
source-address = 0.0.0.0
destination-address-mask = 0.0.0.0
destination-address = 0.0.0.0
Src-Port-Cmp = none
source-port = 0
Dst-Port-Cmp = none
dest-port = 0
tcp-estab = no
icmp-type = 8
```

```
-----
```

```
Forward = yes
Type = Generic Filter
offset = 0
len = 0
more = no
comp-neq = no
dummyForPadding = 0
mask = 00:00:00:00:00:00:00:00:00:00:00:00
value = 00:00:00:00:00:00:00:00:00:00:00:00
```

## Stopping Echo-Request packets from being forwarded upstream

In the following example, a filter is defined to drop outbound ICMP Echo-Request packets, and the filter is applied to a controller Gigabit Ethernet interface.

- 1 Define the filter to prevent outbound ICMP Echo-Requests.  

```
admin> new filter out-no-echo-req
admin> set output-filters 1 valid-entry = yes
admin> set output-filters 1 Type = ip-filter
admin> set output-filters 1 ip-filter protocol = 1
admin> set output-filters 1 ip-filter icmp-type = echo-request
admin> set output-filters 2 valid-entry = yes
admin> set output-filters 2 forward = yes
admin> write -f
```
- 2 Apply the filter to the Ethernet interface.  

```
admin> read ethernet { 1 8 2 }
admin> set filter-name = out-no-echo-req
admin> write -f
```

Echo-Request packets will no longer be forwarded on the Ethernet interface. All other traffic, including Echo-Replies, will be forwarded.

## Blocking Echo-Requests to a specific IP address

In the following example, a filter is defined to drop inbound ICMP Echo-Request packets to a specific IP address, and the filter is applied to a DSL user connections.

- 1 Define the filter to prevent ICMP Echo-Requests to a specific address.  

```
admin> new filter in-no-ping-sysaddr
admin> set input-filters 1 valid-entry = yes
admin> set input-filters 1 Type = ip-filter
admin> set input-filters 1 ip-filter protocol = 1
admin> set input-filters 1 ip-filter icmp-type = echo-request
admin> set input-filters 1 ip-filter dest-address-mask = 255.255.255.0
admin> set input-filters 1 ip-filter dest-address = 135.112.119.96
admin> set input-filters 2 valid-entry = yes
admin> set input-filters 2 forward = yes
admin> write -f
```
- 2 Configure an IPoA connection for a DSL user and apply the filter.  

```
admin> new connection ipoa-11-40
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 10.40.10.1/24
admin> set session-options data-filter = in-no-ping-sysaddr
admin> set atm-options nailed-group = 540
```



```
admin> write -f
3 Configure a BIR connection for a DSL user and apply the same filter.
admin> new connection bir-11-37
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 10.10.10.0/24
admin> set ip-options local-address = 10.10.10.3/32
admin> set session-options data-filter = in-no-ping-sysaddr
admin> set atm-options nailed-group = 537
admin> set bir-options enable = yes
admin> set bir-options proxy-arp = yes
admin> write -f
```

## Defining route filters

Route filters examine RIP update packets and exclude specific routes from the local system's routing table, or include routes in the table only after modifying their metrics. For route filters, the forwarding action in the filter has no effect.

### Overview of route filter settings

In a FILTER profile, the route-filter subprofile contains the following parameters. The parameters are shown with their default values for input filters. The same values apply for output filter rules.

```
[in FILTER/"" ]
filter-name* = ""

[in FILTER/"" :input-filters[1]]
valid-entry = no
forward = no
type = gen-filter

[in FILTER:input-filters[1]:route-filter]
source-address-mask = 0.0.0.0
source-address = 0.0.0.0
route-mask = 0.0.0.0
route-address = 0.0.0.0
add-metric = 0
action = none
```

If you set the parameters in an input filter, only inbound packets are examined. If you set them in an output filter, only outbound packets are examined.

| Parameter   | Setting                                                                                                        |
|-------------|----------------------------------------------------------------------------------------------------------------|
| filter-name | Name of the FILTER profile, up to 36 characters. You apply a filter to an interface by referring to this name. |
| valid-entry | Enable/disable the rule.                                                                                       |
| forward     | This setting does not apply to route filters.                                                                  |

| Parameter           | Setting                                                                                                                                                                                                                                                                                                                                  |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| type                | Type of filter (ip-filter, gen-filter, route-filter, and ethernet-filter). Only the parameters in the corresponding subprofile are applicable for the rule.                                                                                                                                                                              |
| source-address-mask | Mask to be applied to the source-address value before comparing that value to the source address of a RIP update packet.                                                                                                                                                                                                                 |
| source-address      | IP address. After applying the source-address-mask value, the system compares the result to the source address in a RIP packet.                                                                                                                                                                                                          |
| route-mask          | Mask to be applied to the destination address of a route.                                                                                                                                                                                                                                                                                |
| route-address       | IP address. After applying the route-mask value, the system compares the result to routes in a RIP packet. If it finds a route with a matching destination, it takes the specified action.                                                                                                                                               |
| add-metric          | Number from 1 to 15, to be added to the metric value for a route that matches the filter rule, if the specified value for the action parameter is Add.                                                                                                                                                                                   |
| action              | An action to take on a route that matches the filter rule. Valid values are none (the default), accept (accept the route by allowing it to affect the routing table), deny (deny the route by not allowing it to affect the routing table), or add (add the value of the add-metric parameter to the route metric and accept the route). |

## Sample route filters

The following sample filters show how to exclude a route and how to modify a route's metrics.

### Sample route filter that excludes a route

In this example, the defined input filters accept all inbound RIP packets except those with a destination of 90.0.0.0. Following are the commands entered to define the filter, and the system's responses:

```
admin> new filter route-test
admin> set input 1 valid = yes
admin> set input 1 type = route-filter
admin> set input 1 route route-mask = 255.0.0.0
admin> set input 1 route route-address = 90.0.0.0
admin> set input 1 route action = deny
admin> set input 2 valid = yes
admin> set input 2 type = route-filter
admin> set input 2 route action = accept
admin> write -f
```

In this sample route filter, any route that matches rule 1 is rejected, and all other routes are accepted (because they match rule 2).

### Sample route filter that configures a route's metric

In this example, an output filter identifies the route 11.0.0.0 in outbound RIP packets and assigns a high metric to that route.

```
admin> new filter metrics
admin> set output 1 valid = yes
admin> set output 1 type = route-filter
admin> set output 1 route route-mask = 255.0.0.0
admin> set output 1 route route-address = 11.0.0.0
admin> set output 1 route add-metric = 7
admin> set output 1 route action = add
admin> write -f
```

## Defining Ethernet input filters

The system supports Ethernet input filters for filtering input streams on the basis of Ethernet type and MAC addresses. Ethernet filters examine the EtherType field in the Ethernet header and forward or drop only the specified type of traffic (for example, PPPoE traffic). MAC address filters examine source and destination MAC addresses in the Ethernet header and forward or drop frames from a particular source or destination.

Currently, Ethernet filters apply only to input streams.

### Overview of ethernet filter settings

Following are the parameters of the new ethernet-filter subprofile, shown with default settings:

```
[in FILTER/"""]
filter-name* = ""
[in FILTER/""":input-filters[1]]
valid-entry = no
forward = no
type = gen-filter
[in FILTER/""":input-filters[1]:ethernet-filter]
source-mac-address = 00:00:00:00:00:00
destination-mac-address = 00:00:00:00:00:00
ethernet-type = 00:00
```

| Parameter   | Setting                                                                                                        |
|-------------|----------------------------------------------------------------------------------------------------------------|
| filter-name | Name of the FILTER profile, up to 36 characters. You apply a filter to an interface by referring to this name. |
| valid-entry | Enable/disable the rule.                                                                                       |

| Parameter               | Setting                                                                                                                                                                                                                                                                                |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| forward                 | Forwarding action for the rule. The default value of no causes the system to discard matching packets.                                                                                                                                                                                 |
| type                    | Type of filter (ip-filter, gen-filter, route-filter, and ethernet-filter). Only the parameters in the corresponding subprofile are applicable for the rule.                                                                                                                            |
| source-mac-address      | A source MAC address. A six-byte hexadecimal number uniquely representing the source host; for example, 12:34:56:78:9a:bc. The system will forward or discard packets from this address, depending on how the rule is defined.                                                         |
| destination-mac-address | A destination MAC address. A six-byte hexadecimal number uniquely representing the destination host; for example, 12:34:56:78:9a:bc. The system will forward or discard packets addressed to this host, depending on how the rule is defined.                                          |
| ethernet-type           | Ethernet type. A two-byte hexadecimal number representing the Ethernet protocol. PPPoE values are 8863 (PPPoE Discovery Stage packets), and 8864 (PPP Session Stage packets). The system will forward or discard packets of the specified types, depending on how the rule is defined. |

### **Sample PPPoE and MAC address filter**

The following sample filter shows both PPPoE filtering (explicitly allowing only inbound PPPoE traffic) and MAC address filtering (explicitly discarding packets to and from specific devices).

```
admin> new filter
admin> set filter-name = enet-filter2
```

The following commands configure input-filter 1 to forward packets with ethertype = 0x8863 (PPPoE Discovery Stage) and input-filter 2 to forward packets with ethertype = 0x8864 (PPP Session Stage). All other, nonmatching types of Ethernet frames will be discarded.

```
admin> set input 1 Type = ethernet-filter
admin> set input 1 valid-entry = yes
admin> set input 1 forward = yes
admin> set input 1 ethernet-filter ethernet-type = 88:63
admin> set input 2 Type = ethernet-filter
admin> set input 2 valid-entry = yes
admin> set input 2 forward = yes
admin> set input 2 ethernet-filter ethernet-type = 88:64
```

The next commands configure input-filter 3 to discard packets from the MAC address 11:22:33:44:55:66 and to the MAC address 66:22:33:44:55:11:

```
admin> set input 3 Type = ethernet-filter
admin> set input 3 valid-entry = yes
admin> set input 3 ethernet-filter source-mac = 11:22:33:44:55:66
admin> set input 3 ethernet-filter destination-mac = 66:22:33:44:55:11
admin> write -f
```

## ***Applying a filter to an interface***

This section describes how to apply a filter to an interface.

### ***Profile settings for applying a filter***

To apply a filter to an interface, set the following parameters (shown with default settings):

```
[in CONNECTION/"" :session-options]
data-filter = ""

[in CONNECTION/"" :ip-options]
route-filter = ""

[in IP-INTERFACE/{ { shelf-1 first-control-module 2 } N}]
route-filter = ""

[in ETHERNET/{ shelf-1 first-control-module 2 }]
filter-name= ""

[in VLAN-ETHERNET/{ { shelf-1 first-control-module 2 } N }]
filter-name = ""
```

| Parameter    | Setting                                                                                                                                                                                                                                                     |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| data-filter  | Name of a FILTER profile containing IP or Ethernet filter rules to be applied to the WAN interface.                                                                                                                                                         |
| route-filter | Name of a FILTER profile containing route filter rules to be applied to the interface.                                                                                                                                                                      |
| filter-name  | Name of a FILTER profile containing IP or Ethernet filter rules to be applied to the Ethernet or VLAN interface. For all types of bridged VLANs, the configured filter must be an Ethernet input filter. For routed VLANs, the filter must be an IP filter. |

### ***Applying filters to a CPE interface***

When you apply an IP filter in a CPE CONNECTION profile, it prevents certain inbound packets from reaching the LAN side of the system, or certain outbound packets from reaching the CPE router. When you apply a route filter, it affects which routes received in RIP updates on this interface will be allowed to affect the Stinger system's routing table.

## Filter Configuration

### Applying a filter to an interface

---

Following is an example of applying both an IP filter and a route filter to a terminating PVC:

```
admin> read connection cpe-1
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 3.3.3.200/30
admin> set atm-options vpi = 8
admin> set atm-options vci = 100
admin> set session-options data-filter = ip-filter1
admin> set ip-options route-filter = route-filter1
admin> set atm-options nailed-group = 201
admin> write -f
```

### Applying a filter to an Ethernet interface

A filter applied to an Ethernet interface takes effect immediately. If you change any settings in a FILTER profile, the changes apply as soon as you save the FILTER profile. The filter you apply affects only non-VLAN traffic on the interface. Non-VLAN traffic consists of untagged frames and frames with VLAN IDs that do not match configured VLAN-ETHERNET profiles.



**Note** Use caution when applying a filter to the Ethernet interface. You could inadvertently render the Stinger IP DSLAM inaccessible from the LAN.

The following set of commands applies an Ethernet filter to the IP DSLAM Gigabit Ethernet interface:

```
admin> read ethernet { 1 8 2 }
admin> set filter-name = enet-filter1
admin> write -f
```

### Applying a filter to a VLAN Ethernet interface

A filter you apply to a VLAN Ethernet interface affects only traffic routed or bridged to the VLAN. The following set of commands creates a bridged VLAN interface and applies an Ethernet filter to it:

```
admin> new vlan-ethernet
admin> set interface-address physical shelf = shelf-1
admin> set interface-address physical slot = first-control-module
admin> set interface-address physical item-number = 2
admin> set interface-address logical-item = 50
admin> set enabled = yes
admin> set bridging-options bridging-group = 451
admin> set bridging-options bridge = yes
```

```
admin> set bridging-options bridge-type = transparent-bridging
admin> set filter-name = enet-filter2
admin> write -f
```

For an example of applying a filter to a routed VLAN interface, see “Applying an IP filter to a routed VLAN” on page 3-38.

## ***Applying an Ethernet filter to a VLAN bridging connection***

The following commands create a subscriber VLAN bridging connection and apply the sample Ethernet filter (enet-filter2) created in “Sample PPPoE and MAC address filter” on page 11-20.

```
admin> new connection
admin> set station = vlan-pppoe-1
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options ip-routing-enabled = no
admin> set bridging-options bridging-group = 451
admin> set bridging-options bridge = yes
admin> set bridging-options bridge-type = transparent-bridging
admin> set session-options data-filter = enet-filter2
admin> set atm-options vci = 60
admin> set atm-options nailed-group = 125
admin> write -f
```

## ***Administrative tools for filters***

If you have access to the debug environment, you can use the debug-level info np command to verify that packets are being discarded and forwarded properly on a filtered interface. For details about using this command and enabling the debug environment, see “Network processor-related diagnostics” on page A-20.

The system also supports the system-level filterdisp command for displaying information about filters in use on all terminating connections.

```
admin> filterdisp
ID Username   Src   Route-Filter  Data-Filter  Call-Filter  TOS-Filter
-----
015  baj         loc                f1
021  man         loc                f1
022  man2        loc                f1
023  man3        loc
<end user list> 4 active user(s)
```

The first column of the output displays a session ID number, followed by a username and the name of the filter. To display details for a particular session, specify the session ID as an argument on the filterdisp command line.

## Filter Configuration

### Administrative tools for filters

---

For example, the following sample output shows that no filters are applied to session 23:

```
admin> filterdisp 23
Hostname:      man3
No associated filters
```

The following sample output shows filters applied to an externally authenticated session:

```
admin> filterdisp 17
Hostname:      edleung
searching for external filters...
Externally obtained filters exist

Data Filter
Direction: Out

Forward = yes
Type = IP Filter
protocol = 0
source-address-mask = 0.0.0.255
source-address = 1.1.1.2
destination-address-mask = 0.0.0.0
destination-address = 0.0.0.0
Src-Port-Cmp = none
source-port = 0
Dst-Port-Cmp = none
dest-port = 0
tcp-estab = no
```

Internally, the system assigns each configured filter an identifier, and the filter ID is attached to a connection. The IDs are allocated based on the sequence in which the filters are created, starting with 1. The following command displays all filters and their corresponding IDs, ordered alphabetically by filter name:

```
admin> filterdisp -a
Name          Id
-----
baj           2
man           1
man2          3
man3          4
Total number of entries found 4
```

The following command displays the filter name for a specified filter ID:

```
admin> filterdisp -n 3
Filter name for filter id 3 is man2
```

The next command returns the filter ID for the specified filter name:

```
admin> filterdisp -i man3
Filter Id for filter man3 is 4
```



---

# IP Control Module Diagnostics

# A

|                                                        |      |
|--------------------------------------------------------|------|
| Enabling the debug environment .....                   | A-2  |
| Gigabit Ethernet diagnostics.....                      | A-2  |
| IGMP diagnostics .....                                 | A-4  |
| PIM-SM diagnostics .....                               | A-10 |
| VLAN-related diagnostics.....                          | A-13 |
| SAR-related diagnostics .....                          | A-19 |
| Network processor-related diagnostics.....             | A-20 |
| Multicast-capable high-bandwidth LIM diagnostics ..... | A-27 |
| SNMP MIB for GMAC and VLAN statistics .....            | A-29 |
| PIMv2 MIB support .....                                | A-35 |

This appendix describes the `gmac`, `igmp`, and `pim` system-level commands. Information provided by the `gmac` command is also accessible to an external management utility through the `ip2kstats.mib` management information base (MIB). Information displayed by the `pim` command can also be accessed through the PIMv2 MIB (currently defined in `draft-ietf-pim-mib-v2-01.txt`) under `experimental 61`.

In addition, this appendix describes the following debug-level commands, which are not documented in the *Stinger Reference* and are not supported.

- `diag igmpsp`
- `diag igmp`
- `brtbls`
- `ifmgr`
- `diag brtbl`
- `sar`
- `info np`



**Caution** Debug-level commands can be used to display low-level details about IP DSLAM operations. However, they are introduced into the system for development purposes, and are not part of the supported software environment. Use debug-level commands with caution!

## Enabling the debug environment

To access the debug environment, log in using the default **super** (super-user) profile. For example:

```
admin> auth super
Password:
```

The default password for this account is **Ascend**. If the password has not been changed, you should change it now to prevent unauthorized super-user access to the system.

The debug environment contains many hidden commands and parameters that are not intended for general use. The debug environment is not supported, and the documentation provided here is not comprehensive.



**Caution** Under most circumstances, debug commands are not required for monitoring Stinger IP DSLAM operations, and under some circumstances, these commands might produce undesirable results. Use the information with caution. Contact Alcatel-Lucent OnLine Customer Support at <http://www.lucent.com/support> with questions or concerns.

## Gigabit Ethernet diagnostics

The **gmac** command provides diagnostic output about the Gigabit Ethernet media access controller (GMAC) driver.

### **gmac**

**Description** Provide diagnostics on the Gigabit Ethernet (GigE) driver.

**Permission level** system

**Usage** **gmac** [options]

| Command element | Description                                                                                                            |
|-----------------|------------------------------------------------------------------------------------------------------------------------|
| -v              | Show <b>gmac</b> version.                                                                                              |
| -i [-u/d]       | With no additional option, initialize/reset the GigE port.<br>-i -u Force GigE link up.<br>-i -d Force GigE link down. |
| -n              | Set up network processor to communicate with GMAC.                                                                     |
| -s              | Set up a SAR channel for communicating with GMAC.                                                                      |

| Command element | Description                                                                                                                                                                                              |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -l [-i/e/d/p]   | Loopback.<br>-l -i   Set port for internal loopback.<br>-l -e   Set port for external loopback.<br>-l -d   Set port for no loopback<br>-l -p   Run loopback test for Ethernet power-on self test (POST). |
| -p              | Ping test.                                                                                                                                                                                               |
| -r              | Read a PHY register.                                                                                                                                                                                     |
| -w              | Write to a PHY register.                                                                                                                                                                                 |
| -d [-c/a/e]     | With no additional option, display all statistics.<br>-d -c   Clear GMAC statistics.<br>-d -a   Display ATM statistics.<br>-d -e   Display Ethernet statistics.                                          |
| -t              | Set debug level (0 through 3).                                                                                                                                                                           |
| -?              | Display a summary of commands.                                                                                                                                                                           |

**Example** The `gmac -n` command sets up the network processor for communication with the GMAC port:

```
super> gmac -n
NP setup for gmac done.
```

**Example** The `gmac -s` command sets up a SAR channel for communicating with the GMAC port:

```
super> gmac -s
GMAC: SAR conn. open with vpi = 0, vci = 200
```

**Example** The `gmac -v` command displays the GMAC version:

```
super> gmac -v
GMAC version : 0x0b
```

**Example** The `gmac -i` command resets the Gigabit Ethernet port:

```
super> gmac -i
gigE port reset.
```

**Example** The `gmac -d` command displays the total transmit and receive statistics for the GigE interface of the IP control module. For details about the `gmac -d` output fields, see the descriptions of MIB objects in “Total transmit statistics” on page A-30 and “Total receive statistics” on page A-32. For example, the following command displays the current GMAC statistics:

```
super> gmac -d
Gigabit Ethernet port statistics :

txOctetsLow      = 1040
txOctetsHigh     = 0
txGoodPackets    = 4
```

```

txPkt64           = 0
txPkt65127        = 0
txPkt128255       = 0
txPkt256511       = 4
txPkt5121023      = 0
txPkt1024Max      = 0
txPktDefer        = 0
txPktUndSz        = 0
txUnderFlow       = 0
txPfcb = 0
txPfcb = 0
txRfcb = 0
txRfcc = 0
txOverflow        = 0
txAlmostFull      = 0

rxOctetsLow       = 1646718
rxOctetsHigh      = 0
rxGoodPackets     = 2059
rxPkt64           = 766
rxPkt65127        = 0
rx128255          = 0
rx256511          = 160
rx5121023         = 0
rx1024Max         = 1133
rxMacType         = 0
rxCrcErrors       = 0
rxUnderSize       = 0
rxOverSize        = 0
rxAlmostFull      = 0
rxOverRun         = 0
rxMulticastPackets = 1896
rxBroadcastPackets = 46
rxJabber          = 0
rxPfc = 0
rxRfc = 0

```

## ***IGMP diagnostics***

The **igmp** command provides information about IGMP multicast operations. The **profile** argument has been added to support MCAST-SERVICE profiles.

The **diag igmpsp** and **diag igmp** commands are available only in the debug environment. See “Enabling the debug environment” on page A-2 for related information.

### **igmp**

**Description** Displays or clears multicast information about IGMP groups and clients.

**Permission level** system

The system-level **igmp** command supports new set of arguments for displaying information obtained by IGMP snooping. The **igmp** command is supported both on the shelf and the individual LIM slots. Following is the new usage statement, which shows both the existing and new arguments:

```
admin> igmp
```

Igmp Commands Usage:

```
igmp <mbone | groups | clients | slots | vc >
    <profile [mcast-profile] | delete [grp_addr][if_num]>
```

Igmp Snooping Commands Usage:

```
igmp <smbone | sgroups [bridgeNum] | sclients [bridgeNum]>
    <sslots | sdelete [bridgeNum][grp_addr][if_num]>
```

| Command element         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |              |       |        |       |  |  |           |    |     |            |    |     |            |    |     |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|-------|--------|-------|--|--|-----------|----|-----|------------|----|-----|------------|----|-----|
| mbone                   | Display multicast backbones. For example:<br>admin> igmp mbone<br>Mbone is currently:<br>Slot 1:8 ifNum = 1                                                                                                                                                                                                                                                                                                                                                                                                                                |              |       |        |       |  |  |           |    |     |            |    |     |            |    |     |
| groups                  | Display currently registered multicast group addresses and interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |              |       |        |       |  |  |           |    |     |            |    |     |            |    |     |
| clients                 | Display multicast clients.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |              |       |        |       |  |  |           |    |     |            |    |     |            |    |     |
| slots                   | Display multicast enable slots.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |              |       |        |       |  |  |           |    |     |            |    |     |            |    |     |
| vc                      | Check the multicast VCs that are configured and active at any point of time. For example:<br>admin> igmp vc<br>Current VCs:<br><table><tr><th>Profile Name</th><th>IfNum</th><th>Active</th></tr><tr><td colspan="3">-----</td></tr><tr><td>oc3-multi</td><td>17</td><td>yes</td></tr><tr><td>oc3-multi2</td><td>25</td><td>yes</td></tr><tr><td>oc3-multi1</td><td>24</td><td>yes</td></tr></table>                                                                                                                                       | Profile Name | IfNum | Active | ----- |  |  | oc3-multi | 17 | yes | oc3-multi2 | 25 | yes | oc3-multi1 | 24 | yes |
| Profile Name            | IfNum                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Active       |       |        |       |  |  |           |    |     |            |    |     |            |    |     |
| -----                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |              |       |        |       |  |  |           |    |     |            |    |     |            |    |     |
| oc3-multi               | 17                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | yes          |       |        |       |  |  |           |    |     |            |    |     |            |    |     |
| oc3-multi2              | 25                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | yes          |       |        |       |  |  |           |    |     |            |    |     |            |    |     |
| oc3-multi1              | 24                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | yes          |       |        |       |  |  |           |    |     |            |    |     |            |    |     |
| profile [mcast-profile] | Display multicast profiles. If you specify an MCAST-SERVICE profile name, the command also displays the CONNECTION profiles associated with that profile. This option is available only on the system control module.<br><br>For example, the following command displays information about an MCAST-SERVICE profile named test:<br>admin> igmp profile test<br>Mcast Service : test<br>Filter Type : MCAST_FILTER_INCLUSIVE<br>Filter List :<br>239.10.10.10 - 239.10.10.10 ( Count = 1 )<br>Connection Profiles :<br>1) stg61<br>2) stg63 |              |       |        |       |  |  |           |    |     |            |    |     |            |    |     |

| Command element               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| delete [grp_addr<br>[if_num]] | <p>With the delete option alone, delete all currently registered multicast groups and their members. If a group address is specified, delete all members of that group. If a group address and interface number are specified, delete that member of the specified group. For example:</p> <pre>admin&gt; igmp delete 226.1.1.2 12</pre> <p>LOG notice, Shelf 1, Controller-1, Time: 10:25:53--<br/>Multicast client 226.1.1.2 link DOWN interface<br/>number 12</p> <p>When executed from the shelf controller prompt, only groups that have been joined from the trunk side can be deleted. You must open a session to the LIM to delete groups that have been joined from the LIM side.</p> <p>sdelete [bridgeNum]      Delete all the members of this<br/>[grp_addr]                  group, for this bridge on this<br/>                                 slot.</p> <p>sdelete [bridgeNum]      Delete a unique instance.<br/>[grp_addr][if_num]</p> |

**Example** The `igmp client` command shows local MBONE and IGMP client interfaces. For example:

```
super> igmp client
IGMP Clients
```

| Client   | Version | RecvCount | CLU | ALU |
|----------|---------|-----------|-----|-----|
| 2(Mbone) | 2       | 0         | 0   | 0   |
| 4(Mbone) | 2       | 0         | 0   | 0   |
| 1(Mbone) | 2       | 3370      | 0   | 0   |
| 5        | 2       | 0         | 0   | 0   |
| 6        | 2       | 0         | 0   | 0   |

The output contains the following fields:

| Field     | Description                                                                                                                                                                                                                                                                                                                                                                             |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Client    | In <code>igmp client</code> command output, the <code>Client</code> field displays the interface ID (the <code>ifIndex</code> value) on which the client resides. The 1 value represents the Gigabit Ethernet interface. Other numbers are WAN interfaces, numbered according to when they became active. The interfaces labeled (Mbone) receive multicast data from multicast routers. |
| Version   | IGMP version.                                                                                                                                                                                                                                                                                                                                                                           |
| RecvCount | Number of IGMP messages received on that interface.                                                                                                                                                                                                                                                                                                                                     |
| CLU/ALU   | CLU is current line utilization, and ALU is average line utilization. Both indicate the percentage of bandwidth used across this interface. If bandwidth utilization is high, some IGMP packet types are not forwarded.                                                                                                                                                                 |

**Example** The `igmp groups` command displays information about MBONE interfaces. Details about client member interfaces are maintained on the LIM itself. For example, the following command is invoked on the IP control module of a Stinger system with multicast clients on a DSL interface in slot 6 and an MBONE configured on a Gigabit Ethernet interface:

```
super> igmp groups
IGMP Group address Routing Table
Up Time: 0d 1:25:05
Group Address  MemberIf  Expire time  Counts
239.100.100.5  *          (Mbone)      0::0 S2
                Slot 1:6
```

The next commands open a session with the DSL LIM and invoke `igmp groups` on the LIM:

```
super> open 1 6
dads1-atm-24-1/6> igmp groups
IGMP Group address Routing Table
Up Time: 0d 1:23:46
Group Address  MemberIf  Expire time  Counts
239.100.100.5  6         00:04:07     0::0 S2
```

When the command is executed on the LIM, the output displays details about the corresponding client member interface. The output contains the following fields:

| Field         | Description                                                                                                                                                                                                                                                                                                                                  |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Group address | Multicast address used for the group. An asterisk indicates the IP multicast address being monitored. If a group has no members, the system forwards multicast traffic for the group to the MBONE interface (the default route).                                                                                                             |
| MemberIf      | Interface ID of multicast group members.                                                                                                                                                                                                                                                                                                     |
| Expire time   | When this membership expires. The system sends out IGMP queries every 60 seconds, so the expiration time is usually renewed. If the expiration time is reached, the system removes the entry from the table. If the field contains periods, this membership never expires. A string of periods means that the default route never times out. |
| Counts        | Number of packets forwarded to the client, the number of packets dropped due to lack of resources, and the state of the membership.                                                                                                                                                                                                          |

**Example** The `igmp slot` command displays information about slots supporting IGMP clients:

```
super> igmp slot
IGMP Client Slots
Shelf:Slot      Group      SendCount
1:8             230.0.0.9  0
1:5             230.0.0.9  0
1:2             230.0.0.9  0
```

The output contains the following fields:

| Field      | Description                                     |
|------------|-------------------------------------------------|
| Shelf:Slot | Shelf and slot card the MBONE connection is on. |
| Group      | Interface number of connection.                 |
| SendCount  | Number of packets sent across the interface.    |

**Example** The `igmp profile` command displays information about MCAST-SERVICE profiles. If you execute the command on the control module and specify a profile name, the command will also display the CONNECTION profiles associated with it.

```
super> igmp profile
```

IGMP Service Profiles

```
Service Name      : gold-service
Filter Type       : MCAST_FILTER_INCLUSIVE
Filter List       :
                  224.255.129.120
                  224.225.129.119

Service Name      : bronze-service
Filter Type       : MCAST_FILTER_INCLUSIVE
Filter List       :
                  224.255.129.119
```

The output contains the following fields:

| Field        | Description                                               |
|--------------|-----------------------------------------------------------|
| Service Name | Name of the multicast service profile.                    |
| Filter Type  | Inclusive multicast group filtering in the named profile. |
| Filter List  | Multicast group addresses to be filtered.                 |

## igmp hosts

**Description** Displays host addresses per member interface for fast-leave purposes. This command is not available on the controller. To use it, you must open a session to the module on which the member interfaces are established.

**Permission level** system

**Usage** `igmp hosts`

**Example** The following commands can be used to verify the hosts that have joined a multicast group on an ADSL module in slot 6:

```
admin> open 1 6
```

```
dadsl-atm-24-1/6> igmp hosts
```

| Group Address | MemberIf | Host Address |
|---------------|----------|--------------|
| 230.1.1.1     | 1        | 10.10.10.2   |
|               |          | 10.10.10.3   |



In the command output, 230.1.1.1 is the group address and two hosts, 10.10.10.2 and 10.10.10.3, have joined the group on interface 1.

## diag igmpsp

**Description** Enable low-level diagnostics on IGMP services profiles.

**Permission level** debug

**Usage** diag igmpsp

**Example** This example first uses a system-level command to view IGMP multicast service profile characteristics. It then uses the debug environment to display low-level diagnostic information.

```
super> igmp profile
IGMP Service Profiles
```

```
Service Name      : gold-service
Filter Type       : MCAST_FILTER_INCLUSIVE
Filter List       :
                  239.100.100.4
```

The profile named gold-service has been applied to a connection on the second port of that LIM. The following commands open a session with the ADSL LIM and display diagnostics on the service profiles:

```
super> open 1 6
```

```
dadsl-atm-24-1/6> diag igmpsp
mcastJoinFilterVerify: Received CLASS D address 239.100.100.4 in report
_profileNameCompare: Compare ( gold-ser with gold-ser )
mcastJoinFilterVerify: [ gold-ser ] is multicast service profile for [ term-
6-2 ] connection
```

The next output indicates that the system accepted the message received from a client on the second LIM interface:

```
mcastJoinFilterVerify: Received CLASS D address 239.100.100.4 in report
_profileNameCompare: Compare ( gold-ser with gold-ser )
mcastJoinFilterVerify: [ gold-ser ] is multicast service profile for [ term-
6-2 ] connection
```

## diag igmp

**Description** Enable low-level diagnostics on IGMP protocol messages.

**Permission level** debug

**Usage** diag igmp

**Example** This example displays IGMP join messages from a client. The following output shows a group message to a client on LIM slot 6, port 6 and the client's response requesting to join multicast group 230.0.0.9:

```
super> diag igmp
igmp debug is ON
igmpParseMsg: IGMP packet to 230.0.0.9 type 6 on interface 6 port 6
```

```

Receiving Version 2 Response from 6
Joining Group 230.0.0.9
IGMP: Joining new group 230.0.0.9
    _sendIGMPTableUpdateMsg: sending IGMP_TAB_ADD to 1:8
    _sendUpdateMsgToShelf : client 6 join group 230.0.0.9 vRouterID 0
igmpParseMsg: IGMP packet to 230.0.0.9 type 6 on interface 6 port 6
Receiving Version 2 Response from 6
IGMP: Refreshing group 230.0.0.9 input ifNum 6

```

## PIM-SM diagnostics

The `pim` command provides information about PIM-SM v2 operations.

### pim

**Description** Displays PIM-related information.

**Permission level** system

**Usage** `pim [ group | rp | nbr | if ifnum | bsr | hash group-addr ]`

| Command element              | Description                                                   |
|------------------------------|---------------------------------------------------------------|
| <code>groups</code>          | Displays information about multicast groups.                  |
| <code>rp</code>              | Displays group-RP mappings.                                   |
| <code>nbr</code>             | Displays information about PIM neighbor routers.              |
| <code>if ifnum</code>        | Displays PIM interface statistics.                            |
| <code>bsr</code>             | Displays information about candidate BSRs or the elected BSR. |
| <code>hash group-addr</code> | Displays the best RP for a group or group range.              |

**Example** The `pim groups` command displays information about all multicast groups. For example:

```
super> pim groups
```

```

Group Addr  RP/Source Addr upJPTimer  Tree(Rpt/Spt)
223.1.1.1   192.168.101.1  40          RPT
224.4.4.4   10.10.10.10   70          RPT

```

#### Output field (pim groups) Description

|               |                                                                                                    |
|---------------|----------------------------------------------------------------------------------------------------|
| Group Addr    | Multicast group address for the entry                                                              |
| RP/SourceAddr | With the current software version, the RP/Source address is the RP's IP address.                   |
| upJPTimer     | The interval in seconds following which the Stinger IP DSLAM sends another PIM join/prune message. |
| Tree(Rpt/Spt) | With the current software version, the value of this field is always RPT (RP tree).                |

**Example** The `pim rp` command displays PIM information for candidate Rendezvous Points (C-RPs) for IP multicast groups. When the local router is the BSR, this

information is obtained from received C-RP-Advertisements. When the local router is not the BSR, this information is obtained from received RP-Set messages. For example:

```
super> pim rp
```

| Group        | RP-Address | RPF neighbor | Priority | holdtime |
|--------------|------------|--------------|----------|----------|
| 224.0.0.0/8  | 1.1.1.3    | 1.1.1.3      | 40       | 75:62    |
| 234.0.0.0/8  | 1.1.1.3    | 1.1.1.3      | 40       | 75:62    |
| 234.3.0.0/16 | 1.1.1.3    | 1.1.1.3      | 40       | 75:62    |
| 234.4.0.0/16 | 1.1.1.3    | 1.1.1.3      | 40       | 75:62    |
| 234.5.0.0/16 | 1.1.1.3    | 1.1.1.3      | 40       | 75:62    |
| 234.6.0.0/16 | 1.1.1.3    | 1.1.1.3      | 40       | 75:62    |

Static entries (if any) appear following a Static Entries label at the end of the list. Static entries represent the PIM-GROUP-RP-MAPPING profile entries. Fields in the command output have the following meaning:

| Output field (pim rp) | Description                                                                                                                                                                                                                                                                                                                                                                     |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Group                 | The IP multicast group address and group mask for which this entry contains information about the C-RP.                                                                                                                                                                                                                                                                         |
| RP-Address            | IP address of the C-RP.                                                                                                                                                                                                                                                                                                                                                         |
| RPF neighbor          | IP address of the Reverse Path Forwarding (RPF) neighbor-router to which a join message for this group would be directed to under certain circumstances. The RPF neighbor for an RP is calculated when the Stinger IP DSLAM receives an IGMP join message for a group range that uses the RP.                                                                                   |
| Priority              | Priority of this mapping, to be used in selecting a C-RP if multiple mappings are found for a group range.                                                                                                                                                                                                                                                                      |
| holdtime              | Holdtime of the C-RP. If the Stinger is the BSR, the first value in this field represents the holdtime received in the C-RP message for this RP and the second value shows the time left before this C-RP will be removed. If the Stinger is not the elected BSR, the first value in this field represents the holdtime received in BSR messages and second value will be zero. |

**Example** The `pim nbr` command displays information about PIM neighbors on active PIM interfaces in the Stinger IP DSLAM. For example:

```
super> pim nbr
```

| Neighbor  | Interface | Priority | Holdtime | DR  |
|-----------|-----------|----------|----------|-----|
| 1.1.101.2 | 2         | 2        | 100      | YES |

Fields in the command output have the following meaning:

| Output field (pim nbr) | Description                                                                          |
|------------------------|--------------------------------------------------------------------------------------|
| Neighbor               | The IP address of the PIM neighbor.                                                  |
| Interface              | The value of <code>ifIndex</code> for the interface used to reach this PIM neighbor. |

| Output field (pim nbr) | Description                                                                                                                                                                                         |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Priority               | Hello priority of the neighbor. The 0 value indicates that the neighbor does not support the priority option, or the neighbor supports the priority option but has an assigned hello priority of 0. |
| Holdtime               | Interval in seconds before the Stinger times out the neighbor if no hello message is received.                                                                                                      |
| DR                     | Whether the neighbor is DR on the interface.                                                                                                                                                        |

**Example** The `pim if` command displays information and statistics about the specified PIM interface. For example:

```
super> pim if 1
pimHelloIntvl      30
pimHelloHoldtime   105
pimHelloPriority    1
pimJpIntvl         60
pimJpHoldtime      210
pimLanPruneDelay   5000
pimOdDelay         2500
pimDR              FALSE
genId              22305411
PIM Statistics
    4 packets received
    0 bad checksum packets received
    0 bad version packets received
    3 hello packet received
    0 join/prune packets received
    0 Boot strap packets received
    1 C-RP Adv packets received
    5 packets transmitted
    4 hello packets sent
    0 join/prune packets sent
    1 boot strap packets sent
```

The statistics list the number of various packet types sent or received on this interface. They reach a maximum value and are then reset to zero. The `pimDR` field indicates whether the Stinger IP DSLAM is DR on the interface. The other fields in `pim if` command output correspond to the `pim-options` configuration for the interface, as described in “PIM options in the IP-INTERFACE and CONNECTION profiles” on page 10-7.

**Example** The `pim bsr` command shows information about candidate BSRs within the domain if the Stinger is acting as candidate BSR, otherwise it shows information about the elected BSR. For example, the following output occurs when the Stinger IP DSLAM has been elected BSR:

```
super> pim bsr
Stinger BSR State : ELECTED_BSR
Details of CURRENT BSR:
BSR IP Address      : 1.1.1.101
BSR Interface       : 1
BSR Priority         : 255
```

```
BSR holdtime      : 57
BSR Current Frag Tag : 717
BSR HASH masklen   : 30
```

Fields in the command output have the following meaning:

| Output field (pim bsr) | Description                                                                                                                                                              |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Stinger BSR State      | State of the system relevant to BSR election. If C-BSR is not enabled in the IP-GLOBAL profile, this field displays <i>Stinger is not a C-BSR</i> .                      |
| BSR IP Address         | The IP address of the bootstrap router (BSR) for the local PIM region.                                                                                                   |
| BSR Interface          | The value of <i>ifIndex</i> for the interface used to reach the BSR.                                                                                                     |
| BSR Priority           | Priority of the current BSR, from 0 to 255.                                                                                                                              |
| BSR holdtime           | The bootstrap holdtime when the BSR is a C-RP in the local domain. If <i>Stinger</i> is elected BSR, holdtime represents the interval after which next BSM will be sent. |
| BSR Current Frag Tag   | The value used in the current BSM to identify fragmented BSMs.                                                                                                           |
| BSR HASH masklen       | A value (30 by default) used to calculate the hash value for a group range when two RPs have the same priority.                                                          |

**Example** The `pim hash` command displays the IP address of the best RP for a group or group range. For example:

```
super> pim hash 234.1.1.1
Best RP for group 234.1.1.1 is 1.1.1.10
```

To determine the best RP for a particular group, the system searches the dynamic list of RPs first. If the system does not find an RP in the dynamic list, the *Stinger* system looks into the static RP list. For dynamic entries, the system selects the best RP on the basis of the longest prefix match, RP priority, and the RP IP address. For static entries, the selection of best RP is based on longest prefix match.

## VLAN-related diagnostics

The `brtbls` (bridge tables), `ifmgr`, `diag brtbls`, and `vlanstats` commands are available only in the debug environment. See “Enabling the debug environment” on page A-2 for related information.

### brtbls

**Description** The `brtbls` command supports diagnostics related to VLAN bridge circuits.

**Permission level** debug

**Usage** brtbcls [-c] | [-i n] [-p ifnum] [-s n] [-r]

| Command element | Description                                                                    |
|-----------------|--------------------------------------------------------------------------------|
| -c              | Show all bridge circuits.                                                      |
| -i n            | Show interfaces on bridge circuit <i>n</i> , or within bridge-group <i>n</i> . |
| -p ifnum        | Show partner information on <i>ifnum</i> .                                     |
| -s n            | Show stack user VLANs on bridge circuit <i>n</i> .                             |

**Example** This example displays bridge table information for VLAN bridge circuits.

```
super> brtbcls -c
bridgeGroup  bridgeType  circuit  interfaceList  items on List
      10         VLAN_CKT  0x84019f30      8400f040         1
     1000        VLAN_STACK 0x84016300      82c4d830         3
2 circuits.
```

**Example** The next command displays interfaces on bridge circuit 11:

```
super> brtbcls -i 11
ifNum  ifType  ifName
   11    135  vlan11
   13    49   wan13
```

The next command lists the interfaces in bridge group 1000:

```
super> brtbcls -i 1000
ifNum  ifType  ifName
   23    49   wan23
   22    49   wan22
   20   135  vlan1000
```

**Example** The next command shows paired bridge-circuit interface information for interface 13:

```
super> brtbcls -p 13
[ifNum 13 iff 80c4aa88] <---> [ifNum 11 iff 80c4a678]
```

**Example** The `brtbcls -r` command shows the “router” interface (the default uplink path) for a bridge group. For a bridge group on STP-enabled ports, it displays two “router” interfaces. For example:

```
super> brtbcls -r 10
IF-NUM  PhyAddr
-----  -
   24  {1 8 3}
   23  {1 8 2}
```

## **ifmgr**

**Description** The `ifmgr` command displays interface-table entries, toggles the debug display, and marks an interface as enabled or disabled. To see the VLAN interface related information, this command should only be entered from the Control Module.

**Permission level** debug

**Usage** ifmgr [options]

| Command element          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| [-r vrouter]             | Display routing entries. If a virtual router name is specified on the command line, the command displays only the table of the virtual router. If no virtual router name is specified, the command displays the tables for all virtual routers.                                                                                                                                                                                                                                                                                                                                     |
| [-r vrouter] -k ifNum    | Enables manual deletion of a virtual interface on a BIR connection. If ifNum is the interface number of a virtual interface on a BIR connection, the command deletes that virtual interface.<br><br>If ifNum is the interface number of the main configured BIR connection, the command deletes all virtual interfaces associated with that interface. (The main BIR connection interface is never deleted through the use of this command.)<br><br>If ifNum is neither the interface number of a virtual interface or a main BIR interface, the command displays an error message. |
| -d [ifname ifnum]        | Displays interface table entries.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| -n conn-profile name     | Display slot and interface number of the connection.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| [up down] [ifnum ifname] | Enables or disables the specified interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

**Example** The ifmgr -d command displays the interface table:

```
super> ifmgr -d
bif slot sif u m p ifname      host-name  remote-addr  local-addr
-----
000 1:08 000 *    ie0        -          0.0.0.0/32    9.9.9.55/32
001 1:08 001 *    ie1        -          0.0.0.0/32    0.0.0.0/32
002 1:08 002 *    ie2        -          0.0.0.0/32    10.10.10.2/32
032 1:08 007 *    ie-stp-1-21-  -          0.0.0.0/32    0.0.0.0/32
033 1:08 008 *    ie1-2111    -          0.0.0.0/32    31.31.31.2/32
034 1:08 009 *    ie2-2111    -          0.0.0.0/32    31.31.31.2/32
008 1:08 008 *    lo0        -          0.0.0.0/32    127.0.0.1/32
009 0:00 000 *    rj0        -          0.0.0.0/32    127.0.0.2/32
010 0:00 000 *    bh0        -          0.0.0.0/32    127.0.0.3/32
011 1:08 000 *    wanabe     -          0.0.0.0/32    127.0.0.3/32
012 0:00 000 *    local     -          0.0.0.0/32    127.0.0.1/32
013 0:00 000 *    mcast     -          0.0.0.0/32    224.0.0.0/32
014 0:00 000 -    tunnel0    -          0.0.0.0/32    127.0.0.5/32
015 0:00 000 *    vr0_main   -          0.0.0.0/32    127.0.0.4/32
017 0:00 000 *    sar1-61-17 -          0.0.0.0/32    0.0.0.0/32
018 0:00 000 *    sar1-62-18 -          0.0.0.0/32    0.0.0.0/32
019 0:00 000 *    sar1-9-19  -          0.0.0.0/32    0.0.0.0/32
020 0:00 000 *    sar1-8-20  -          0.0.0.0/32    0.0.0.0/32
021 0:00 000 *    sar1-9-21  -          0.0.0.0/32    0.0.0.0/32
022 0:00 000 -    sip0      -          0.0.0.0/32    0.0.0.0/32
023 1:03 001 *    p wan23    mpoa-102   5.5.5.5/32    5.5.5.1/32
```

```

024 0:00 000 *   wan24      -          0.0.0.0/32      0.0.0.0/32
025 0:00 000 *   wan25      -          0.0.0.0/32      0.0.0.0/32
026 1:08 007 *   ie-lag-1   -          0.0.0.0/32      10.10.1.83/32
027 1:08 008 *   ie-lag-1-1 -          0.0.0.0/32      10.10.7.83/32
<end>

```

The `ifmgr -d` command now indicates a virtual STP interface. For example, in the following output, virtual interface 32 is designated `ie-stp-1-2111` and the real interface 33 (a routed VLAN) is designated `ie1-2111`.

```
super> ifmgr -d
```

| bif | slot | sif | u | m | p | ifname       | host-name | remote-addr | local-addr    |
|-----|------|-----|---|---|---|--------------|-----------|-------------|---------------|
| 000 | 1:08 | 000 | * |   |   | ie0          | -         | 0.0.0.0/32  | 9.9.9.55/32   |
| 001 | 1:08 | 001 | * |   |   | ie1          | -         | 0.0.0.0/32  | 0.0.0.0/32    |
| 002 | 1:08 | 002 | * |   |   | ie2          | -         | 0.0.0.0/32  | 10.10.10.2/32 |
| 032 | 1:08 | 007 | * |   |   | ie-stp-1-21- |           | 0.0.0.0/32  | 0.0.0.0/32    |
| 033 | 1:08 | 008 | * |   |   | ie1-2111     | -         | 0.0.0.0/32  | 31.31.31.2/32 |
| 034 | 1:08 | 009 | * |   |   | ie2-2111     | -         | 0.0.0.0/32  | 31.31.31.2/32 |

#### Command element

#### Description

`bif`

Bundle interface number. There is one interface number per bundle, including Multilink Protocol (MP) connections. This number is the global interface-table number.

`slot`

Shelf and slot to which the interface is assigned.

`sif`

Slot interface.

`u`

Whether the interface is enabled (\*) or disabled (-).

`m`

The interface is part of an MP bundle.

`p`

Whether the interface is permanent. A `p` indicates a permanent interface. A hyphen (-) or a blank indicates that the interface is not permanent. A permanent interface is an interface that is configured in the command-line interface and stored in Stinger NVRAM. All the Ethernet interfaces and the interfaces based on CONNECTION profiles are permanent. Transient interfaces are those the Stinger unit builds from RADIUS. These interfaces have no interface entry when the connection is not active.



| Command element | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ifname          | <p>Interface name. Many of the interface names are described in “Displaying the interface table” on page 4-3.</p> <p>The names <code>ie1</code> and <code>ie2</code> represent GigE interfaces.</p> <p>The names <code>sar&lt;shelf&gt;-&lt;slot&gt;-&lt;ifnum&gt;</code> represent SAR interfaces.</p> <p>The names <code>ie-lag-&lt;grp&gt;</code> represent LAG main interfaces, where <code>&lt;grp&gt;</code> is the index of the ETHER-GROUP profile. LAG virtual interfaces are shown the following name:</p> <p><code>ie-lag-&lt;grp&gt;-&lt;num&gt;</code></p> <p>where <code>&lt;grp&gt;</code> is the index of the ETHER-GROUP profile and <code>&lt;num&gt;</code> is the last number of the IP-INTERFACE profile corresponding to the virtual interface.</p> |
| host-name       | Hostname of remote device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| remote-addr     | Remote address of device as configured in a CONNECTION profile.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| local-addr      | Address of the local interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

**Example** The next command displays information about the Gigabit Ethernet interface:

```
super> ifmgr -d 1
iff                0x82ec86cc
inUse:             Yes
hostName:
dialoutName:
Authentication Source: In: local      Out: local
ExternFilters:    No
ExternRoutes @    0
miscInfo @        0
reDirectDest:     0.0.0.0
DLCI routeId:     0
MP(P) id:         0
rtIf: 1:08:1
virtual id: 0, virtual next @ -1, virtual main @ -1
minor device:     1
device status:    0x221
output func:      0x801b5f18
input func:       0x801b61a0
mtu:              1500
ip_addr:          201.168.53.123
dstip_addr:       0.0.0.0
netmask:          255.255.255.0
net:              201.168.53.0
subnet:           201.168.53.0
bcast:            201.168.53.255
nbcast:           201.168.53.255
```

```

directed-bcast:      yes
management only:    no
macaddr:            00c07b65d579
inp_qcnt:           0
out_qcnt:           0
nexthop:            0.0.0.0
proxy_arp_mode:     0
proxy_arp_head:     0
vRouterID:          0
if_redirServer:     0.0.0.0
if_redirPort:       0
if_redirPort:       0
ATMP tunnel:        DISABLED
No associated connection profile
SNMP ifType:        6
multicastServiceProfile :
multicastMaxGroups  :      0

```

## diag brtbls

**Description** Enable diagnostic printf input for bridge tables.

**Permission level** debug

**Usage** diag brtbls

**Example** The next commands enable and disable bridge table diagnostics.

```

super> diag brtbls
brtbls                               ( Bridge Tables diagnostic )
brtbls debug is ON
super> diag brtbls
brtbls                               ( Bridge Tables diagnostic )
brtbls debug is OFF

```

## vlanstats

**Description** Display or clear VLAN statistics. For details about the vlanstats output fields, see the descriptions of MIB objects in “VLAN statistics tables” on page A-34.

**Permission level** debug

**Usage** vlanstats [-c] {shelf slot item} vlan-id

| Command element | Description                              |
|-----------------|------------------------------------------|
| -c              | Clear statistics for the specified VLAN. |

**Example** The following command displays statistics about VLAN 1 on the first controller's GigE interface:

```

super> vlanstats { 1 8 2 } 1
rxOctetsHigh      : 0
rxOctetsLow       : 0
rxFrames          : 0

```

```
rxUnicastFrames : 0
rxMulticastFrames : 0
rxBroadcastFrames : 0

txOctetsHigh : 0
txOctetsLow : 0
txFrames : 0
txUnicastFrames : 0
txMulticastFrames : 0
txBroadcastFrames : 0
```

**Example** The following command clears the statistics for VLAN 1:

```
super> vlanstats -c { 1 8 2 } 1
Statistics for VLAN 1 cleared
```

## ***SAR-related diagnostics***

The sar command is available only in the debug environment. See “Enabling the debug environment” on page A-2 for related information.

### **sar**

**Description** Display the Segmentation and Reassembly (SAR) interface and routing tables, protocol statistics, and active sockets.

**Permission level** debug

**Usage** sar [options]

| <b>Command element</b> | <b>Description</b>                                  |
|------------------------|-----------------------------------------------------|
| -b                     | Show free transmission buffer list.                 |
| -c                     | Clear SAR statistics.                               |
| -i                     | Reset the SAR (destructive).                        |
| -l                     | List open channels.                                 |
| -p                     | Configure transmission packet display.              |
| -q                     | Configure receive packet display.                   |
| -m                     | Show SAR control memory information.                |
| -r                     | Read shaper.                                        |
| -s                     | Show SAR statistics.                                |
| -t                     | Execute SAR PHY test (control module).              |
| -v                     | Show SAR virtual circuit table.                     |
| -w                     | Write shaper.                                       |
| -x                     | List open connections.                              |
| -y                     | Loopback cell statistics (shelf).                   |
| -z [-d -o -i -c -a]    | List open WAN connections (Stinger control module). |

**Example** The following command analyzes contents of a frame on Gigabit Ethernet to check the VLAN ID:

```
super> sar -p -a -100
SAR: now dumping the contents of all transmitted packets

super> ping -c 1 20.1.2.10
PING 20.1.2.10 (20.1.2.10): 56 data bytes
tx 1/61(d) @ a1d46b00 packet len 102
TX packet: (task "_brouterPacketTask" at 0x81697c40, time: 8907.26) 102
octets @ 0xa1d46b00
[0000]: 01 20 01 02 03 04 00 d0 52 01 02 04 81 00 00 64 . . . . .
R. . . . .d
[0010]: 08 00 45 00 00 54 c9 2a 00 00 ff 01 c6 6e 14 01 ..E..T.*
. . . . .n.
[0020]: 02 04 14 01 02 0a 08 00 8a 47 23 a5 00 00 60 00 . . . . .
.G#...`.
[0030]: ea 12 00 00 00 00 00 00 00 00 00 00 00 00 00 00 . . . . .
. . . . .
[0040]: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 . . . . .
. . . . .
[0050]: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 . . . . .
. . . . .
[0060]: 00 00 00 00 . . . .
```

## Network processor-related diagnostics

The `info np` command is available only in the debug environment. See “Enabling the debug environment” on page A-2 for related information.

### info np

**Description** Display hardware, protocol, and connection information.

**Permission level** debug

**Usage** `info np option`

| Option                           | Description                                                                                                                                                                                                                                 |
|----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>addr</code>                | Display PayloadPlus <code>addr</code> information                                                                                                                                                                                           |
| <code>arp [vrouter]</code>       | Display ARP entries. If a virtual router name is specified on the command line, the command displays only the table of the virtual router. If no virtual router name is specified, the command displays the tables for all virtual routers. |
| <code>bridge option</code>       | <code>group group-num</code> Display bridging group details.<br><code>table group-num ifnum</code> Display bridging table details.                                                                                                          |
| <code>conn [slot] [ifnum]</code> | Display connections.                                                                                                                                                                                                                        |
| <code>did [didnum]</code>        | Display destination ID (DID) table contents.                                                                                                                                                                                                |
| <code>ipqos [ifnum]</code>       | Display IP QoS details.                                                                                                                                                                                                                     |

| Option                       | Description                                                                                                                                                                                                                                     |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| lacp                         | Display link aggregation group (LAG) details.                                                                                                                                                                                                   |
| lns                          | Display LNS entries                                                                                                                                                                                                                             |
| mb1                          | Display Mobile entries                                                                                                                                                                                                                          |
| mgrp                         | Display multicast groups.                                                                                                                                                                                                                       |
| mpm                          | Display the VLAN-MAC to port mapping table in the network processor.                                                                                                                                                                            |
| pm bwdist                    | Display IP2100 port-manager bandwidth distribution.                                                                                                                                                                                             |
| pm rates                     | Display IP2100 port-manager bandwidth rates.                                                                                                                                                                                                    |
| psl lport [cnt]<br>[rspslot] | Display lport-scheduler information                                                                                                                                                                                                             |
| qid qidnum [rspslot]         | Display Queue information                                                                                                                                                                                                                       |
| rsc                          | Display the resources collection list.                                                                                                                                                                                                          |
| route [vrouter]              | Display routing entries. If a virtual router name is specified on the command line, the command displays only the table of the virtual router. If no virtual router name is specified, the command displays the tables for all virtual routers. |
| sed didnum [cnt]             | Display SED parameters indexed by DID                                                                                                                                                                                                           |
| sid sidnum [rspslot]         | Display scheduler information                                                                                                                                                                                                                   |
| spoof [ifnum]                | Display spoof details.                                                                                                                                                                                                                          |

**Example** Without arguments, the `info np conn` command displays all connection handles. For example:

```
super> info np conn
Connection Handles
if slot act cPort dPort type uctl udata qos dctl ddata qos
19 2 Y 000001 000001 LIM 0/124 0/123 default 0/124 0/123 default
1 8 Y 000000 000000 GE 0/864 0/864 default 0/864 0/864 default
6 8 Y 000000 000000 GE 0/864 0/864 default 0/864 0/864 default
7 8 Y 000000 000000 GE 0/864 0/864 default 0/864 0/864 default
24 8 Y 000000 000000 GE 0/864 0/864 default 0/864 0/864 default
```

To dump all the connection handles for a given slot, or for a given port on a slot, specify the slot/port numbers on the command line. For example:

```
super> info np conn 1 2 22
Connection Handles
if sh/sl act cPort dPort type uctl udata dctl ddata
22 1/2 Y 000000 000001 ATM 0/124 0/60 0/60 0/124
Encaps: AAL5-LLC - MTU=1560 (50)
```

Control & Data Service Queues:

```
-----
ControlQID          354   Lport          0001
Scheduler            UBR   CosQueue          1

DataQID             352   Lport          0010
```

Scheduler UBR CosQueue 1

PDU Ids:

-----

| PduID  | port   | vpi/vci | Tree |
|--------|--------|---------|------|
| 000156 | 0x0401 | 0/124   | 3098 |
| 000157 | 0x0200 | 0/60    | 3098 |
| 000158 | 0x0401 | 2/124   | 3098 |
| 000159 | 0x0200 | 0/60    | 3098 |

VLAN Configuration : Vlan is not enabled

Packet Flow configuration: No Packet Flow is attached

Multicast Configuration : Multicast is not enabled

Bridge Configuration : Bridging is not enabled

Filter Configuration : No Filters applied

QOS Configuration :

|            |           |
|------------|-----------|
| Downstream | : default |
| Upstream   | : default |

DID List:

| Did    | Type    | QId | cosId |
|--------|---------|-----|-------|
| 572252 | UNICAST | 353 | 1     |
| 572253 | UNICAST | 355 | 1     |

OAM DID List:

| Did    | Type    |
|--------|---------|
| 572254 | UNICAST |
| 572255 | UNICAST |
| 572256 | UNICAST |
| 572257 | UNICAST |

If the connection is a LAG interface, the command shows an additional “cloned” data service queue below the Control & Data Service Queues section of the output. This represents a duplicate queue for the second Gigabit Ethernet port, which is used in the distribution of frames.

For example:

```
super> info np conn 1 8 29
```

Connection Handles

| if | sh/sl | act | cPort  | dPort  | type | uctl | udata | dctl | ddata |
|----|-------|-----|--------|--------|------|------|-------|------|-------|
| 29 | 1/8   | Y   | 000000 | 000000 | LAG  | 0/0  | 0/0   | 0/0  | 0/0   |

Encaps: ETH/ATM - MTU=1500 (50)

Control & Data Service Queues:

-----

|                 |     |       |        |
|-----------------|-----|-------|--------|
| DataQID         | 320 | Lport | 0x0020 |
| DataQID (clone) | 324 | Lport | 0x0021 |

```
Scheduler                      UBR   CosQueue                      3

PDU Ids:
-----
None.

VLAN Configuration             : Vlan-Id 7

Packet Flow configuration: No Packet Flow is attached

Multicast Configuration       : Multicast is not enabled

Bridge Configuration          : Bridge Group 20      Transparent Bridging enabled

Filter Configuration          : No Filters applied

QOS Configuration              :
    Downstream                 : default
    Upstream                   : default
```

VLAN CONNECTION: next DIDs are for vlan-id 7

Principal DID List:

| Did     | Type      | QId | cosId | Cloned |
|---------|-----------|-----|-------|--------|
| 50512   | MULTICAST | 322 | 2     | No     |
| 50513   | MULTICAST | 326 | 2     | Yes    |
| 1048558 | UNICAST   | 323 | 3     | Yes    |

Layer 2 DID List: Empty

**Example** The `info np mpm` command displays the VLAN-MAC to port mapping table in the network processor. For example:

```
super> info np mpm
MAC Address  VlanID IfNum  Did  CnHdl
-----
00:00:E2:19:BD:16  2111 33 572243 8894a240
```

The first column of the output displays the MAC address of the gateway. The second column displays the VLAN ID associated with the MAC address in the first column. The third column displays the interface number of the Gigabit Ethernet port on which the MAC frame was received. The last two columns display the DID used to send a frame to the gateway and its associated connection handle.

**Example** To display all routes on the network processor, use the `info np route` command. For example:

```
super> info np route
Local If Routes for vrouter:main
Destination      Gateway          IF          DID      needArp
```

```

1.1.1.3/32      -          local      000000    N
9.0.0.0/8       -          ie1-1     000000    N
9.9.9.91/32     -          local      000000    N
9.9.9.92/32     -          local      000000    N
12.0.0.0/8      -          ie1        000000    N
12.12.12.12/32  -          local      000000    N
127.0.0.0/8     -          bh0        000000    N
127.0.0.1/32    -          local      000000    N
127.0.0.2/32    -          rj0        000000    N
135.254.196.0/24 210.210.210.1 ie0        000000    N
210.0.0.0/8     -          ie0        000000    N
210.210.210.75/32 -        local      000000    N
224.0.0.0/4     -          mcast     000000    N
224.0.0.1/32    -          local      000000    N
224.0.0.2/32    -          local      000000    N
224.0.0.9/32    -          local      000000    N
224.0.0.13/32   -          local      000000    N
255.255.255.255/32 -        ie0        000000    N
Route Table for vrouter:main
Destination      Gateway          IF              DID      needArp
1.1.1.0/24       1.1.1.12        wan19           000000    Y
1.1.1.12/32      1.1.1.12        wan19           000000    Y
Total 20 routes <end>

```

**Example** The info np lacp command displays LAG information, such as whether the constituent ports are distributing or not. For example:

```

super> info np lacp
802.3ad Link Aggregation status on NP:
-----

```

```

o LAG 1:
  - Ge port 1: distributing
  - Ge port 2: distributing

* Ge port 1: group=1, LAG-IfNum=32
* Ge port 2: group=1, LAG-IfNum=32

```

**Example** The info np mgrp displays multicast group information. For example:

```

super> info np mgrp
IP Address  MaxMtu  Member  FirstDid  LastDid  LstInList  NdDrain
224.1.1.1   1610    1       12855     12855    12855      No

```

**Example** To display IP QoS details, IP QoS monitoring must be enabled on the interface. For information about enabling QoS monitoring, see “Administrative tools for monitoring IP QoS” on page 5-47. The info np ipqos command displays information being monitored on the specified interface. For example:

```

super> info np ipqos 19
IP QOS monitoring information for IF 19
IP QOS profile: Pav-Pbit
Number of priority      :4          Lowest priority        :1
Default classification :provided    Total Mcast Rules      :0
Classified DIDs         :3          Classification rules    :4

```



Mapping priority/COS

P[0]=x P[1]=5 P[2]=x P[3]=x P[4]=4 P[5]=3 P[6]=2 P[7]=x

Main DID list

-----

| DID   | Qid | COS | Type  |
|-------|-----|-----|-------|
| 43514 | 285 | 5   | ucast |

Classified DID list

-----

| DID   | Qid | COS | Type  |
|-------|-----|-----|-------|
| 43516 | 282 | 2   | ucast |
| 43518 | 283 | 3   | ucast |
| 43520 | 284 | 4   | ucast |

**Example** To display the IP2100 NP port-manager fixed-rate (guaranteed) rate allocation based on NP-PORT profile settings, use the `info np pm bwdist` command. For example:

```
super> info np pm bwdist
```

NP Port Manager Rates

NP PM Bandwidth Distribution

System Aggregate Allocated BW : 432581

Aggregate User Requested BW : 433416

Available system BW : 3486580

| Port     | Requested | Allocated | Status      |
|----------|-----------|-----------|-------------|
| {1 1 0}  | 36805     | 36805     | BW SERVICED |
| {1 2 0}  | 9359      | 9393      | BW SERVICED |
| {1 3 0}  | 36805     | 36805     | BW SERVICED |
| {1 4 0}  | 9359      | 9393      | BW SERVICED |
| {1 5 0}  | 36805     | 36805     | BW SERVICED |
| {1 6 0}  | 9359      | 9393      | BW SERVICED |
| {1 7 0}  | 9359      | 9393      | BW SERVICED |
| {1 10 0} | 9359      | 9393      | BW SERVICED |
| {1 11 0} | 9359      | 9393      | BW SERVICED |
| {1 12 0} | 36805     | 36805     | BW SERVICED |
| {1 13 0} | 9359      | 9393      | BW SERVICED |
| {1 14 0} | 9359      | 9393      | BW SERVICED |
| {1 15 0} | 36805     | 36805     | BW SERVICED |
| {1 16 0} | 9359      | 9393      | BW SERVICED |
| {1 17 1} | 2540      | 2555      | BW SERVICED |
| {1 17 2} | 2540      | 2555      | BW SERVICED |
| {1 18 1} | 2540      | 2555      | BW SERVICED |
| {1 18 2} | 2540      | 2555      | BW SERVICED |
| {1 0 2}  | 62500     | 62654     | BW SERVICED |
| {1 0 3}  | 62500     | 62654     | BW SERVICED |

**Example** To display the IP2100 NP port-manager nonfixed-rate (nonguaranteed) rate allocation, use the `info np pm rates` command. For example:

```
super> info np pm rates
```

NP Port Manager Rates

```

Port{1 1 0} NonFixedRate :552083
Port{1 2 0} NonFixedRate :140896
Port{1 3 0} NonFixedRate :552083
Port{1 4 0} NonFixedRate :140896
Port{1 5 0} NonFixedRate :552083
Port{1 6 0} NonFixedRate :140896
Port{1 7 0} NonFixedRate :140896
CM#8 NonFixedRate :10004
CM#9 NonFixedRate :10004
Port{1 10 0} NonFixedRate :140896
Port{1 11 0} NonFixedRate :140896
Port{1 12 0} NonFixedRate :552083
Port{1 13 0} NonFixedRate :140896
Port{1 14 0} NonFixedRate :140896
Port{1 15 0} NonFixedRate :552083
Port{1 16 0} NonFixedRate :140896
Port{1 17 1} NonFixedRate :38339
Port{1 17 2} NonFixedRate :38339
Port{1 18 1} NonFixedRate :38339
Port{1 18 2} NonFixedRate :38339
Port{1 0 2} NonFixedRate :939814
Port{1 0 3} NonFixedRate :939814
PCI NonFixedRate :8483

```

**Example** To display detailed information about a bridge group, use the `info np brig` command. Its argument is the internal bridge group number. For example:

```

super> info np brig 1
Number of interfaces in bridge group #50 : 3
Internal bridge group number           : 1
Time left to Mac Aging Timer expiry    : 30

rif type bport UDID  FDID  host-name
-----
0020 wan  0001  43526 12600 minal-3.28
0021 wan  0002  43528 12601 minal-13.4
0022 vlan 0003  43530 12602 vlan22

Flood list pattern :
Type Learn Tree Act  Value  Pattern
BRI   Y   3120  4  00303138  16 48 00000001 00000000

Flood list Information :
list size num  first  last  lastinlist  Not-drained
-----
0    255    3 12600 12602 12602          0

Multicast Groups :
GrpIp      Msize Maxmtu Member First  Last LstInList nDrain
225.1.1.1  255    0    1 13110 13110 13110    0
    Type Learn Tree Act Value  Pattern
    BMP   Y   3123  0 00103336 48 0 00000001 E1010101
235.1.1.1  255    0    1 12855 12855 12855    0

```

```
Type Learn Tree Act Value  Pattern
BMP   Y   3123  0 00103237 48  0 00000001 EB010101
```

**Example** The `info np bridge` command displays bridge table entries. Its arguments are the bridge group number and (optionally) an interface number. For example:

```
super> info np bridge 50
```

| Group | PortBlock | if   | Did   | Destination       |
|-------|-----------|------|-------|-------------------|
| 00050 | ROUTER    | 0022 | 43530 | 00:00:00:00:00:09 |
| 00050 | ROUTER    | 0022 | 43530 | 00:00:00:00:00:0a |
| 00050 | ROUTER    | 0022 | 43530 | 00:00:00:00:00:0b |
| 00050 | ROUTER    | 0022 | 43530 | 00:00:00:00:00:0c |
| 00050 | ROUTER    | 0022 | 43530 | 00:00:00:00:00:0d |

Table has 5 entries.

## Multicast-capable high-bandwidth LIM diagnostics

The `info` command is available only in the debug environment, and only in the LIM interface. See “Enabling the debug environment” on page A-2 for related information.

To use this command, you must first open an administrative session to the module. For example, the following commands open a session to a module in slot 15, and then exit the session:

```
super> open 1 15
ep-hs-gs-48-1/15> quit
```

### info

**Description** On multicast-capable HB LIMs, the `info` command displays information relevant to distributed multicast operations on the LIM. This section describes its use only in the context of multicast-capable high-bandwidth (HB) LIMs. For background information, see “Distributed multicast with HB LIMs (IP2100 only)” on page 9-25.

**Permission level** debug

**Usage** `info <info-type> [port/conn-id/slot] [ifnum]`  
           <info-type> : `ffwbridge mcast`

| Option                                           | Description                                                                                                                               |
|--------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| <code>ffwbridge -e</code>                        | Display all entries in the internal AAL5 bridge. For details about the internal bridge, see “Downstream traffic processing” on page 9-28. |
| <code>ffwbridge -e [&lt;port&gt; vpi vci]</code> | Display the bridge entry for the specified connection.                                                                                    |
|                                                  | <code>&lt;port&gt;</code> The DSL line (starting at 1)                                                                                    |
|                                                  | <code>vpi</code> The VPI on the DSL side.                                                                                                 |

| Option                   | Description                                                                                          |
|--------------------------|------------------------------------------------------------------------------------------------------|
|                          | vci                      The VCI on the DSL side.                                                    |
| ffwbridge -f             | Display the FFW interface. For details about FFW interfaces, see “Conceptual overview” on page 9-26. |
| mcast -p -g              | Display all multicast groups in IGMP proxy (routing) mode.                                           |
| mcast -s -g              | Display all multicast groups in IGMP snooping (bridging) mode.                                       |
| mcast -p -g <ip-address> | Display properties of the specified group in IGMP proxy (routing) mode.                              |
| mcast -s -g <ip-address> | Display properties of the specified group in IGMP snooping (bridging) mode.                          |
| mcast -p -m <ip-address> | Display the member list of the specified group in IGMP proxy (routing) mode.                         |
| mcast -s -m <ip-address> | Display the member list of the specified group in IGMP snooping (bridging) mode.                     |

**Example** For example, the following command displays all bridge entries (all terminated, routed connections) on the HB LIM:

```
ep-hs-gs-48-1/15> info ffwBridge -e
DSL port    VPI    VCI    Bridged to Sonet port    VPI    VCI
      4        0    35              0        0    150
```

**Example** The following command displays more detailed information about a connection on port 4 using VPI 0 and VCI 25:

```
ep-hs-gs-48-1/15> info ffwBridge -e 4 0 35
DSL side
Port                : 4
VPI                  : 0
VCI                  : 35
IW port tag          : 15
IW Rx channel tag    : 16
IW Tx channel tag    : 17
IW Tx flow aggregation tag : 18
Sonet side
Port                : 0
VPI                  : 0
VCI                  : 150
IW port tag          : 19
IW Rx channel tag    : 20
IW Tx channel tag    : 21
IW Tx flow aggregation tag : 22
```

**Example** The following command displays information related to the FFW interface (the interface on which multicast traffic is sourced on the HB LIM).

```
ep-hs-gs-48-1/15> info ffwBridge -f
FFW System
Host IW Rx channel tag : 0
Host IW Tx channel tag : 0
```

Host Tx aggregation tag : 1  
FFW interface  
IW port tag : 6  
IW Rx channel tag : 7  
IW Tx channel tag : 8

**Example** The following command displays information about all multicast groups in proxy mode:

```
ep-hs-gs-48-1/15> info mcast -p -g
```

| GROUP address | Aggr Tag | Nb members |
|---------------|----------|------------|
| 224.1.1.2     | 25       | 1          |

## **SNMP MIB for GMAC and VLAN statistics**

The `ip2kstats.mib` MIB gathers statistics about the GMAC interface of the IP control module, and also collects statistics on a per-VLAN basis. It is implemented as the following proprietary Alcatel-Lucent enterprise MIB:

`ip2kStatsGroup` OBJECT IDENTIFIER ::= { ascend 51 }

The transmit and receive statistics represented in this MIB are also accessible in the command-line interface by using the `gmac -d` command. The VLAN statistics are also accessible in the command-line interface by using the debug-level `vlanstats` command.

### **History maintained at 15-minute intervals**

The system maintains a history of the following fields by averaging them at fixed 15-minute intervals for the last 24 hours (96 intervals):

- Tx and Rx octets (transmit and receive traffic streams)
- Multicast traffic (receive side only)
- Unicast traffic (transmit and receive traffic streams)
- Broadcast traffic (receive side only)
- Cyclic redundancy check (CRC) errors (receive side only)

### **Gigabit Ethernet (GigE) statistics tables**

Gigabit Ethernet statistics are represented in five MIB tables:

- GigE configuration
- GigE interval transmit statistics
- GigE interval receive statistics
- GigE total transmit statistics
- GigE total receive statistics

## Gigabit Ethernet configuration

The `gigEConfigTable` is a configuration table for an IP control module GigE interface. It is indexed by the interface index. This MIB table contains the objects shown in Table A-1:

Table A-1. *GigEConfigTable MIB objects*

| MIB object                      | Description                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>gigEValidIntervals</code> | The number of previous intervals for which data was collected. The value will be 96 unless the interface was brought online within the last 24 hours, in which case the value will be the number of complete 15-minute near-end intervals since the interface has been online. |
| <code>gigELastInitTime</code>   | System time when the GigE port was last initialized. The time is represented in text format.                                                                                                                                                                                   |

## Interval transmit statistics

The `gigETxIntervalTable` is a table containing transmit statistics for the GigE interface in a 15-minute interval. It is indexed by the interface index and the interval number (contained in the `gigETxIntervalNumber` variable). This MIB table contains the objects shown in Table A-2:

Table A-2. *GigETxIntervalTable MIB objects*

| MIB object                                                                    | Description                                                                                                                                                                                   |
|-------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>gigETxIntervalNumber</code>                                             | A number between 1 and 96, where 1 is the most recently completed 15-minute interval and 96 is the 15-minute interval completed 23 hours and 45 minutes prior to interval 1.                  |
| <code>gigETxIntervalOctetsLow</code><br><code>gigETxIntervalOctetsHigh</code> | The lower 32 bits and upper 32 bits of the 64-bit interval transmit packet byte counter, which contains a count of how many bytes were transmitted in error-free packets during the interval. |
| <code>gigETxIntervalUnicastPackets</code>                                     | Total number of Unicast packets transmitted during the interval.                                                                                                                              |

## Total transmit statistics

The `gigETxTotalTable` is a table containing the total transmit statistics for the GigE interface. It is indexed by the interface index. The transmit counters in this table are

also displayed in the output of the `gmac -d` command. See “” on page A-33. This MIB table contains the objects shown in Table A-3:

Table A-3. *GigETxTotalTable MIB objects*

| MIB object              | Description                                                                                                                                                                                                                                                                                                                                  |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| gigETxTotalOctetsLow    | The lower 32 bits and upper 32 bits of the 64-bit transmit packet byte counter, which contains a total count of how many bytes have been transmitted in error-free packets.                                                                                                                                                                  |
| gigETxTotalOctetsHigh   |                                                                                                                                                                                                                                                                                                                                              |
| gigETxTotalGoodPackets  | Total count of packets transmitted without error.                                                                                                                                                                                                                                                                                            |
| gigETxTotalPkt64        | Total count of transmitted 64-byte packets.                                                                                                                                                                                                                                                                                                  |
| gigETxTotalPkt65To127   | Total count of transmitted packets from 65 to 127 bytes in length.                                                                                                                                                                                                                                                                           |
| gigETxTotalPkt128To255  | Total count of transmitted packets from 128 to 255 bytes in length.                                                                                                                                                                                                                                                                          |
| gigETxTotalPkt256To511  | Total count of transmitted packets from 256 to 511 bytes in length.                                                                                                                                                                                                                                                                          |
| gigETxTotalPkt512To1023 | Total count of transmitted packets from 512 to 1023 bytes in length.                                                                                                                                                                                                                                                                         |
| gigETxTotalPkt1024ToMax | Total count of transmitted packets from 1024 bytes up to the MAX bytes in length (1536 for non-jumbo packets and 9728 for jumbo packets). <sup>1</sup>                                                                                                                                                                                       |
| gigETxTotalPktDefer     | Total count of packets that were deferred because the interframe gap was in excess of 96 bits when a packet was available for transmission. The IFG is a delay between code division multiple access/carrier detect (CDMA/CD) packets, intended to provide interframe recovery time for other CSMA/CD sublayers and for the physical medium. |
| gigETxTotalPktUndSz     | Total count of packets less than 64 bytes in length.                                                                                                                                                                                                                                                                                         |
| gigETxTotalUnderFlow    | Total count of packets that were truncated because of an empty FIFO.                                                                                                                                                                                                                                                                         |
| gigETxTotalPfcf         | Total count of Pause Flow Control packets that were sent because the receive FIFO exceeded its highwater mark. The Pause function is a mechanism for full duplex flow control.                                                                                                                                                               |
| gigETxTotalPfcc         | Total count of Pause Flow Control packets that were sent because the client requested them.                                                                                                                                                                                                                                                  |
| gigETxTotalRfcf         | Total count of Reset Flow Control packets that were sent because the receive FIFO went below its lowwater mark.                                                                                                                                                                                                                              |

Table A-3. *GigETxTotalTable MIB objects (Continued)*

| MIB object            | Description                                                                                                           |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------|
| gigETxTotalRfcc       | Total count of Reset Flow Control packets that were sent because the client requested them.                           |
| gigETxTotalOverflow   | Total number of packets in which a write from the physical signaling interface was attempted to a full transmit FIFO. |
| gigETxTotalAlmostFull | Total number of packets in which the transmit FIFO Almost Full flag was set.                                          |

1. Jumbo packets are not currently supported on the Gigabit Ethernet interface of the IP control modules.

## Interval receive statistics

The `gigERxIntervalTable` is a table containing receive statistics for the GigE interface in a 15 minute interval. It is indexed by the interface index and the interval number (contained in the `gigERxIntervalNumber` variable). This MIB table contains the objects shown in Table A-4:

Table A-4. *GigERxIntervalTable MIB objects*

| MIB object                                          | Description                                                                                                                                                                               |
|-----------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| gigERxIntervalNumber                                | A number between 1 and 96, where 1 is the most recently completed 15 minute interval and 96 is the 15 minutes interval completed 23 hours and 45 minutes prior to interval 1.             |
| gigERxIntervalOctetsLow<br>gigERxIntervalOctetsHigh | The lower 32 bits and upper 32 bits of the 64-bit interval receive packet byte counter, which contains a count of how many bytes were received in error free packets during the interval. |
| gigERxIntervalUnicastPackets                        | Count of packets of the Unicast Type received during the interval.                                                                                                                        |
| gigERxIntervalMulticastPackets                      | Count of packets of the Multicast Type received during the interval.                                                                                                                      |
| gigERxIntervalBroadcastPackets                      | Count of packets of the Broadcast Type received during the interval.                                                                                                                      |
| gigERxIntervalCrcErrors                             | Count of packets that failed CRC received during the interval.                                                                                                                            |

## Total receive statistics

The `gigERxTotalTable` is a table containing the total receive statistics for the GigE interface. It is indexed by the interface index. The receive counters in this table are also displayed in the output of the `gmac -d` command. See “” on page A-33. This MIB table contains the objects shown in Table A-5:



Table A-5. GigERxTotalTable MIB objects

| MIB object                                    | Description                                                                                                                                                             |
|-----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| gigERxTotalOctetsLow<br>gigERxTotalOctetsHigh | The lower 32 bits and upper 32 bits of the 64-bit receive packet byte counter, which contains a total count of how many bytes have been received in error free packets. |
| gigERxTotalGoodPackets                        | Total count of packets received without error.                                                                                                                          |
| gigERxTotalPkt64                              | Total count of received 64-byte packets.                                                                                                                                |
| gigERxTotalPkt65To127                         | Total count of received packets from 65 to 127 bytes in length.                                                                                                         |
| gigERxTotalPkt128To255                        | Total count of received packets from 128 to 255 bytes in length.                                                                                                        |
| gigERxTotalPkt256To511                        | Total count of received packets from 256 to 511 bytes in length.                                                                                                        |
| gigERxTotalPkt512To1023                       | Total count of received packets from 512 to 1023 bytes in length.                                                                                                       |
| gigERxTotalPkt1024ToMax                       | Total count of received packets from 1024 bytes up to the MAX bytes in length (1536 for non-jumbo packets and 9728 for jumbo packets). <sup>1</sup>                     |
| gigERxTotalMacType                            | Total count of MAC Type packets received                                                                                                                                |
| gigERxTotalCrcErrors                          | Total count of packets received that failed CRC.                                                                                                                        |
| gigERxTotalUnderSize                          | Total count of packets that were less than 64 bytes in length.                                                                                                          |
| gigERxTotalOverSize                           | Total count of packets received that were greater than the MAX bytes in length (1536 for non-jumbo packets and 9728 for jumbo packets).                                 |
| gigERxTotalAlmostFull                         | Total number of packets in which the receive FIFO exceeded its highwater mark.                                                                                          |
| gigERxTotalOverRun                            | Total number of receive packets in which a write to a full receive FIFO was attempted.                                                                                  |
| gigERxTotalMulticastPackets                   | Total number of receive packets that were of the Multicast Type.                                                                                                        |
| gigERxTotalBroadcastPackets                   | Total number of receive packets that were of the Broadcast Type.                                                                                                        |
| gigERxTotalJabber                             | Total count of receive packets that were classified as Jabber. Jabber is the condition of abnormally long transmissions, usually due to a fault condition.              |
| gigERxTotalPfc                                | Total count of Pause Flow Control packets that were received. The Pause function is a mechanism for full duplex flow control.                                           |

Table A-5. *GigERxTotalTable MIB objects (Continued)*

| MIB object     | Description                                                   |
|----------------|---------------------------------------------------------------|
| gigERxTotalRfc | Total count of Reset Flow Control packets that were received. |

1. Jumbo packets are not currently supported on the Gigabit Ethernet interface of the IP control modules.

## **VLAN statistics tables**

To capture and clear VLAN statistics, the following MIB tables are supported:

- GigE VLAN statistics
- GigE VLAN clear statistics

### **VLAN statistics**

The `GigEVlanStatTable` is the VLAN statistics table for an IP control module GigE interface. It is indexed by the interface index and by VLAN ID (contained in the `gigEVlanId` variable). The counters in this table are also displayed in the output of the `vlanstats debug-level` command. See “`vlanstats`” on page A-18.

This MIB table contains the objects shown in Table A-6:

Table A-6. *GigEVlanStatTable MIB objects*

| MIB object                | Description                                                                                               |
|---------------------------|-----------------------------------------------------------------------------------------------------------|
| gigEVlanId                | VLAN ID. Identifies a specific VLAN on the GigE interface of the IP control module.                       |
| gigEVlanRxOctetsLow       | Receive Frame Byte Counter Low Double Word.                                                               |
| gigEVlanRxOctetsHigh      | Receive Frame Byte Counter High Double Word.                                                              |
| gigEVlanRxGoodFrames      | Receive frame counter. Indicates the total number of frames received on the GigE interface for this VLAN. |
| gigEVlanRxUnicastFrames   | Receive Unicast frame counter. Indicates the total number of unicast frames received for this VLAN.       |
| gigEVlanRxMulticastFrames | Receive Multicast frame counter. Indicates the total number of multicast frames received for this VLAN.   |
| gigEVlanRxBroadcastFrames | Receive Broadcast frame counter. Indicates the total number of broadcast frames received for this VLAN.   |
| gigEVlanTxOctetsLow       | Transmit frame byte counter Low Double word.                                                              |
| gigEVlanTxOctetsHigh      | Transmit Frame Byte Counter High Double word.                                                             |

Table A-6. GigEVLanStatTable MIB objects (Continued)

| MIB object                | Description                                                                                                            |
|---------------------------|------------------------------------------------------------------------------------------------------------------------|
| gigEVLanTxGoodFrames      | Transmit frame counter. Indicates the total number of Ethernet frames transmitted on the GigE interface for this VLAN. |
| gigEVLanTxUnicastFrames   | Unicast transmit frame counter. Indicates the total number of unicast frames transmitted for this VLAN.                |
| gigEVLanTxMulticastFrames | Multicast transmit frame counter. Indicates the total number of multicast frames transmitted for this VLAN.            |
| gigEVLanTxBroadcastFrames | Broadcast transmit frame counter. Indicates the total number of broadcast frames transmitted for this VLAN.            |

## VLAN clear statistics

The GigEVLanClearStatTable is the table for clearing statistics for a given VLAN on the IP control module GigE interface. It is indexed by the interface index and by VLAN ID (contained in the gigEVLanId variable). This table contains the objects shown in Table A-7:

Table A-7. GigEVLanClearStatTable MIB objects

| MIB object           | Description                                                                                                                                                                                  |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| gigEVLanClearStatCmd | Command to clear statistics for a particular VLAN. Values are:<br>none(1)—No action.<br>clearAllStats(1)—Clear all statistics.<br>A GET operation on this object will always return none(1). |

## PIMv2 MIB support

The system provides SNMP MIB support for the PIM protocol as defined in draft-ietf-pim-mib-v2-01.txt (the PIMv2 MIB). The PIMv2 MIB is placed in the MIB tree under experimental 61. The current software supports PIMv2 MIB tables as shown in Table A-8:

Table A-8. Current level of support for PIMv2 MIB tables

| PIMv2 MIB table   | Support in this software version |
|-------------------|----------------------------------|
| pimInterfaceTable | YES                              |
| pimNeighborTable  | YES                              |
| pimIpMRRouteTable | NO                               |
| pimNextHopGroup   | NO                               |

Table A-8. Current level of support for PIMv2 MIB tables (Continued)

| PIMv2 MIB table     | Support in this software version |
|---------------------|----------------------------------|
| pimRPTable          | NO                               |
| pimRPSetTable       | YES                              |
| pimCandidateRPTable | NO                               |
| pimComponentTable   | YES                              |

The `snmpwalk`, `get`, and `getnext` routines are supported for objects in the supported tables. Because Stinger IP DSLAMs operate within a single PIM domain, the component index in `pimComponentTable` is always set to 1. For example, the following `snmpwalk` from an SNMP manager displays information from the `pimInterfaceTable`:

```
$ snmpwalk -m all -O 50.50.50.5 public
experimental.pimMIB.pimMIBObjects.pim.pimInterfaceTable
pimInterfaceAddress.2 = IPAddress: 1.1.1.2
pimInterfaceNetMask.2 = IPAddress: 255.0.0.0
pimInterfaceMode.2 = sparse(2)
pimInterfaceDR.2 = IPAddress: 1.1.1.2
pimInterfaceHelloInterval.2 = 60 seconds
pimInterfaceStatus.2 = active(1)
pimInterfaceJoinPruneInterval.2 = 60 seconds
pimInterfaceTrigHelloInterval.2 = 5 seconds
pimInterfaceHelloHoldtime.2 = 105 seconds
pimInterfaceLanPruneDelay.2 = on(1)
pimInterfacePropagationDelay.2 = 7500 milliseconds
pimInterfaceOverrideInterval.2 = 2500
pimInterfaceGenerationID.2 = on(1)
pimInterfaceJoinPruneHoldtime.2 = 210 seconds
```

The following `snmpwalk` command displays information from the `pimNeighborTable`:

```
$ snmpwalk -m all -O 50.50.50.5 public
experimental.pimMIB.pimMIBObjects.pim.pimNeighborTable
pimNeighborIfIndex.167837953 = 1
pimNeighborExpiryTime.167837953 = Timeticks: (98) 0:00:00.98
pimNeighborMode.167837953 = sparse(2)
pimNeighborLanPruneDelay.167837953 = 0
pimNeighborOverrideInterval.167837953 = 0
pimNeighborTBit.167837953 = off(0)
pimNeighborDRPresent.167837953 = true(1)
```

The following command from an SNMP manager displays information from the `pimRPSetTable`:

```
$ snmpwalk -m all -O s 50.50.50.5 public
experimental.pimMIB.pimMIBObjects.pim.pimRPSetTable
pimRPSetHoldTime.1.224.255.50397441 = 75
pimRPSetHoldTime.1.234.255.50397441 = 75
pimRPSetHoldTime.1.1002.65535.50397441 = 75
pimRPSetHoldTime.1.1258.65535.50397441 = 75
pimRPSetHoldTime.1.1514.65535.50397441 = 75
pimRPSetHoldTime.1.1770.65535.50397441 = 75
pimRPSetExpiryTime.1.224.255.50397441 = Timeticks: (56) 0:00:00.56
```

```
pimRPSetExpiryTime.1.234.255.50397441 = Timeticks: (56) 0:00:00.56
pimRPSetExpiryTime.1.1002.65535.50397441 = Timeticks: (56) 0:00:00.56
pimRPSetExpiryTime.1.1258.65535.50397441 = Timeticks: (56) 0:00:00.56
pimRPSetExpiryTime.1.1514.65535.50397441 = Timeticks: (56) 0:00:00.56
pimRPSetExpiryTime.1.1770.65535.50397441 = Timeticks: (56) 0:00:00.56
```


Following is sample output of an snmpwalk on pimComponentTable:

```
$ snmpwalk -m all -O s 50.50.50.5 public
experimental.pimMIB.pimMIBObjects.pim.pimComponentTable
pimComponentBSRAddress.1 = IPAddress: 1.1.1.101
pimComponentBSRExpiryTime.1 = Timeticks: (45) 0:00:00.45
pimComponentCRPHoldTime.1 = 0 seconds
pimComponentStatus.1 = active(1)
```



---

# Index



## A

### AAL5

- multiplexing options 4-50, 10-8
- PPPoA 4-86
- required encapsulation for PPPoA and PPPoE 4-89
- See also* PPPoA, PPPoE
- terminating PVCs 4-50

address pool definitions, example 4-26

address resolution 4-10

address spoofing

- preventing via IP filters 11-8
- protection for IPoA, BIR, PPPoA, and PPPoE connections 4-47

addresses

- dynamic, requiring acceptance 4-29
- in filters 11-5
- in packet classification 5-10
- slot-based address assignment 4-36
- source interface local addresses 4-42
- virtual routers, effect on 7-2
- WAN virtual interfaces 4-62
- See also* pools

adjacencies, OSPF 8-4

administrative tools

- filters 11-23
- Gigabit Ethernet 2-25
- IGMP A-4
- IP QoS 5-47
- IP routing 4-109
- L2TP 6-26
- OSPF 8-19
- PIM-SM A-10
- VLAN 3-56

### ADSL

- tracking rate adaptation for downstream traffic shaping 5-26

ANSWER-DEFAULTS profile for PPP sessions 4-87

anti-spoofing 4-47

area border router (ABR) capability 8-2

arptable 3-56, 4-109

ASBR. *See* OSPF

ASE preferences, setting 8-17

### ATM QoS

- IP QoS interactions 5-24
- mapping to packet marking value 5-21
- prioritizing IP packet flows based on DSL service contracts 5-29

ATM QoS, guidelines when both multicast and unicast clients are supported 9-21

ATM settings for terminating PVCs 4-49

atm1483type, required settings 4-90

atm-options 4-49

ATM-QOS 1-18

atm-qos-options 5-22

authentication

- Challenge Handshake Authentication Protocol (CHAP) 4-87
- L2TP tunnel authentication 6-3
- OSPF, MD5 (RFC 2178) 8-2
- PPP Authentication Protocol (PAP) 4-87
- server-based tunnel authentication 6-19
- system name used for tunnel authentication 6-4
- tunnel, connection based 6-16

autodetection of PPPoE or PPPoA traffic streams 4-93, 4-96

## B

backbone VLANs 3-27

backup designated router (BDR) 8-3, 8-4

bandwidth 5-23

bandwidth, IP2100 1-9

- configuration 1-10
- fixed-rate and nonfixed-rate 1-9
- NP-PORT profiles 1-9

bidirectional CHAP 4-93

### BIR

- described 4-56
- host route (BIR/32) example 4-58
- interface configuration 4-56

- multiple static address allocation 4-62
- sample configuration on VDSL ports 4-60
- WAN virtual interface configuration 4-62
- WAN virtual interface configuration (DHCP) 4-66
- bir-options 4-56
- bootp-relay 4-68
- BRAS
  - ANSWER-DEFAULTS 4-87
  - bridging IPoA traffic upstream 3-39
  - bridging PPPoA traffic to an upstream BRAS 3-49
  - packet-type 3-41
  - PPPoA configuration 4-91
  - PPPoE configuration 4-95
  - PPPoE service names 4-101
  - recommended call-type setting 4-85
  - required encapsulation types for PPPoA, PPPoE 4-89
  - required setup 4-87
  - terminating traffic 4-89
  - VLAN bridging of IPoA traffic 3-39, 3-50, 3-53
- BRAS profile 3-40
- bras-ip-address 3-40
- bridge groups 3-7
  - defining the upstream interface 3-8
  - DHCP snooping 3-9
  - IGMP snooping 3-8
  - port blocking 3-8
  - PPPoE snooping 3-9
  - timeout for inactive MAC addresses 3-7
- bridged IP routing (BIR). *See* BIR.
- bridged IPoA 3-42
- BRIDGE-GROUP 3-7
- bridging groups 3-6
- bridging-enabled 2-2, 2-11
- bridging-group 3-7
- broadband remote access server. *See* BRAS
- brtbls A-13
- gmac 2-4, A-2
- ifmgr A-14
- igmp 9-10, A-4
- info np A-20
- ipcache 4-109
- iproute 4-109
- netstat 2-4, 4-2, 4-109
- ospf 8-19
- pim A-10
- ping 2-4, 4-109
- sar A-19
- telnet 4-109
- traceroute 4-109
- vlanstats A-18
- which 9-19
- Configuring Gigabit Ethernet interfaces 2-1
- CONNECTION
  - ATM PVCs 4-49
  - BIR 4-56
  - bridging 3-3
  - DHCP relay 4-72
  - IP options 4-50
  - multicast 9-11
  - OSPF 8-13
  - PPP authentication 4-92
  - PPP requirements 4-86
  - PPPoE 4-96
- connection-based L2TP tunnel authentication 6-17
- connections, maximizing the number of 1-8
- connection-specific anti-spoofing 4-47
- costs (OSPF)
  - defaults 8-5
  - defined 8-5
  - parameter, defined 8-11
  - stub areas, and 8-6
- CPE clients, configuring for DHCP address assignment 4-75

## D

## C

- call-type setting for PPP sessions 4-85
- Challenge Handshake Authentication Protocol (CHAP) 4-87
- class boundary addresses, preventing 4-26
- class of service (CoS)
  - overview of defaults 1-17
  - priority queues 1-17
- commands
  - arptable 4-109
  - brtbls A-13
- default routes 4-41
  - protecting from updates 4-13
  - sample configuration 4-41
- default traffic prioritization 1-18
- denial-of-service, protecting against 4-9
- designated router 8-4
- DHCP option 82 4-65
  - DHCP snooping 3-16
  - DHCP snooping configuration 3-18
  - PPPoE snooping 3-20
  - sample configurations 4-75, 4-80, 4-84
  - sending only a hostname upstream 4-71



- suboption configuration 4-70
- suboptions 1 and 2 4-70
- DHCP relay 4-65, 4-68
  - failover to second server 4-68
  - for CPE clients 4-74
  - giaddr 4-69
  - interoperability issues 4-76
  - IP addresses for CPE 4-75
  - LAN issues 4-76
  - per-connection giaddr control 4-72
  - relay agent configuration 4-74
  - sample configurations 4-64, 4-78
  - WAN virtual interfaces on BIR connections 4-66
- DHCP Request packets, snooping 3-14
- DHCP snooping, configuration 3-14
- dhcp-snooping 3-9
- diag brtbls A-18
- diag igmp A-9
- diag igmpsp A-9
- diagnostics A-1
- directed-broadcast-allowed 4-9
- distributed multicast mechanism
  - CPE joins a multicast group 9-27
  - CPE leaves a multicast group 9-28
  - FFW interface multicast behavior 9-27
  - FFW interface scheduling priority, configuring 9-30
  - FFW interfaces 9-27
  - how it works 9-26
- DNS 4-10
- downstream traffic shaping, tracking rate adaptation for 5-26
- DSL bridge CPE 4-58
- DSL Integrated Access Device (IAD) 1-16
- DSL port blocking 3-9
- dynamic routing, RIP 4-7

## E

- ECMP
  - caveat about system-wide load balancing 4-15
  - definition of ECMP routes 4-15
  - flow-based 4-15
  - limitations 4-16
  - packet-based 4-15
  - sample configuration
    - downstream 4-17
    - flow-based downstream 4-19
    - upstream 4-21
- equal cost multipath routes *See* ECMP.
- ETHER-GROUP 2-10

- Ethernet
  - encapsulation on VDSL ports 4-60
  - LACP on IP2100 dual GigE ports 2-9
  - See also* LACP 2-16
- Ethernet egress scheduling and shaping 5-24
- Ethernet input filters. *See* filters.
- Ethernet p-bit marking 5-17
- Ethernet QoS. *See* QoS.
- ethernet-filter 11-19
- ETHER-QOS profile. *See* VDSL.

## F

- fast-leave 9-12
- FFW interfaces. *See* distributed multicast mechanism
- ffw-scheduling-priority 9-30
- FILTER
  - Ethernet input filters 11-19
  - ICMP filters 11-12
  - IP filters 11-2
  - route filters 11-17
  - tcp-estab 11-6
- filters
  - administrative tools 11-23
  - bridged VLAN interface 11-22
  - changing a route's metric 11-19
  - comparison passes in system 11-4
  - explicit default filter rules 11-2
  - ICMP message type 11-13
  - IDs, displaying 11-24
  - IP filter configuration 11-2
  - MAC address filtering 11-20
  - multicast groups and services 9-8
  - multicast groups, exclusive filtering not supported 9-9
  - PPPoE filtering 11-20
  - routes in RIP updates 11-18
  - rule processing 11-2
  - samples for multicast video traffic 9-17
  - specific routes, filtering 11-18
  - TCP-established 11-6
  - VLAN 3-4
  - VLAN interface 3-4
- flow-based ECMP 4-15
- FLOW-SERVICES 3-32
- FPP classification program memory 1-7

## G

- gateway-address 4-40

giaddr field 4-69  
 giaddr, per-connection control of 4-72  
 Gigabit Ethernet  
   administrative tools 2-25  
   configuration options 2-2  
   diagnostics A-2  
   MBONE interface 10-10  
   network processor setup 2-4  
   port redundancy 2-5  
   redundancy 2-5  
   redundancy for RFC 2684 (IPoA) connections 2-6  
   redundancy for VLAN bridging 2-7  
   redundant LAN MBONE 2-8  
   routing terminating PVCs 4-53  
   SAR setup 2-4  
   soft IP interface configuration 2-6  
   statistics A-29  
   verifying packet transfer 2-4  
   verifying the interface setup 2-3  
 global pools, RADIUS 4-25  
 global tunneling configuration 6-10  
 gmac 2-4, A-2  
 group-range-count 9-9

## H

Hello packets 8-10, 10-2  
 hierarchic routing (areas) 8-6  
 high-bandwidth LIMs 1-6  
   changes caused by multicast operations 9-28  
   distributed multicast on 9-25  
 High-bandwidth LIMs. *See distributed multicast mechanism, LIM interfaces.*  
 host route advertisements, suppressing 4-13  
 host routes  
   summarized in advertisements 4-27  
   suppressing advertisement 4-13

## I

ICMP filters 11-12  
 ICMP filters. *See filters.*  
 ifmgr A-14  
 IGMP  
   redundant LAN MBONE 2-8  
 igmp 9-10, A-4  
 igmp hosts A-8  
 IGMP snooping 3-10  
 IGMP timers 9-13  
 igmp-options 9-13

igmp-snooping 3-8  
   4-23, 4-10  
 Integrated Access Device (IAD) 1-16  
 integrated voice and data 1-16  
 interface-group 5-8  
 interfaces, virtual router membership 7-8  
 Internet access 1-16  
 Internet Group Management Protocol (IGMP) *See* multicast  
 IP addresses  
   BIR requirements 4-57  
   DHCP, obtaining for CPE clients 4-75  
   DNS resolution 4-10  
   for Gigabit Ethernet interface 4-7  
   for Gigabit Ethernet redundancy 2-6  
   for terminating PVCs 4-51  
   interface independent (soft) 4-8  
   of next-hop routers 4-40  
   preventing local address spoofing 11-8  
   remote and local (numbered interfaces) 4-52  
   single source address for system 4-9  
   subnet specifications 4-4  
 IP control module  
   model numbers 1-3, 1-6  
   statistics A-29  
 IP control module specifications  
   software capabilities 1-15  
 IP control modules  
   CoS overview 1-17  
   Gigabit Ethernet redundancy 2-5  
   priority queueing 1-18  
   redundant controllers 2-5  
   services overview 1-15  
 IP CoS subsystem *See* QoS  
 IP filters. *See filters*  
 IP interface configuration  
   for GigE redundancy 2-6  
   for soft interface 4-8  
   for virtual interfaces 4-8  
 IP interface table, displaying 7-5  
 IP multicast  
   overview 9-1  
   *See also* multicast 9-2  
 IP multicast. *See distributed multicast mechanism*  
 IP pool chaining  
   defined 4-30  
   local profiles, configuring in 4-31  
   RADIUS, configuring in 4-33  
   *See also* pools  
 IP QoS. *See QoS.*  
 IP routing  
   administrative tools 4-109

- equal cost multipath routes 4-15
  - filtering routing table updates 11-1
  - load balancing 4-15
  - route filters 11-1
  - IP routing table 4-1, 4-2
    - virtual routers, addresses, and 7-2
    - virtual routers, for 7-5
  - IP settings for terminating PVCs 4-50
  - IP2100
    - bandwidth configuration 1-10
    - bandwidth options 1-9
    - features not supported on 1-15
    - introduction to 1-6
    - memory modes 1-7
  - IP2100, bandwidth
    - starvation condition 1-13
  - ipcache 4-109
  - ip-filter 11-3
  - IP-GLOBAL 4-9
    - address pools 4-23
    - DHCP relay settings 4-68
    - multicast forwarding 9-23
    - multicast forwarding 9-4
    - multiple-multicast-server-vc 9-23
    - OSPF 8-9, 8-16
    - RIP options 4-11
    - system address 4-9
  - IP-INTERFACE
    - DHCP relay 4-72
    - OSPF 8-10
    - soft interface 4-8
    - typical settings 4-6
    - virtual interfaces 4-8
  - IPoA bridging to upstream BRAS 3-39
  - ip-options 4-50
  - iproute 4-109
  - IPTV
    - benefits of distributed multicast for 9-26
    - queue size increase for traffic on DSL links 5-26
    - sample configuration 9-19
- ## L
- L2TP
    - call serial number (CSN) 6-10
    - configuration errors causing multiple tunnels 6-5
    - configuring tunnel assignment IDs 6-23
    - connection-based tunnel authentication 6-16
    - enabling operations and authentication 6-6
    - existing tunnels to the same LNS 6-5
    - global options 6-6
    - global tunneling configuration 6-10
      - how client-auth-id settings create parallel tunnels 6-5
      - how the system finds a matching tunnel 6-4
    - links to LNS servers 6-2
    - links to PPP clients 6-2
    - mobile clients 6-11
    - PPP-authenticated connection with two LNS systems 6-14
    - primary and secondary tunnel end points 6-15
    - related network settings 6-3
    - related RADIUS configuration 6-3
    - server-based authentication configuration 6-21
    - server-based tunnel authentication 6-19
    - timers and other variables 6-7
    - tunnel assignment IDs. 6-22
    - tunnel authentication 6-3
    - tunneling, overview 6-2
    - using connection-based tunnel authentication 6-16, 6-17
    - using PPP authentication 6-13
    - using server-based tunnel authentication 6-18
    - using tunnel-server password 6-20
    - using PPP authentication 6-11
  - LAC (L2TP Access Concentrator). *See* L2TP.
  - LACP
    - configuration overview 2-10
    - ETHER-GROUP 2-10, 2-19
    - implementation details 2-16
    - information in command output A-22, A-24
    - IP2100 ports 2-9, 2-17
    - related diagnostics 2-17
    - sample link aggregation configuration 2-14, 2-21
    - virtual ETHERNET profile 2-10, 2-20
  - LAN OSPF interfaces
    - authentication 8-2
    - designated router priority 8-4
  - lan-router-interface-address 3-8
  - layer 2 classification 5-6
  - layer 2 multicasting 3-10
  - layer 3 and interface group classification 5-7
  - LCP keepalives 4-98
  - LIM interfaces
    - multicast clients 9-2, 9-11
    - multicast-capable HB LIM considerations 9-28
    - obtaining an IP address 4-74
    - paired with a VLAN ID 3-5, 3-43, 3-45
    - RFC 2684 PVCs 4-49
  - Limitations
    - IP2000 1-5
    - IP2100 1-15
  - limitations
    - L2TP data sequencing 6-28

L2TP IP QoS, impact 6-28  
 Link Aggregation Control Protocol, *See* LACP  
 link state advertisements. *See* LSAs  
 link-state database  
   adjacencies, and 8-4  
   building 8-7  
   creating 8-7  
   routing table, and 8-7  
   updates 8-4  
 link-state routing algorithm 8-7  
 link-state-enabled 2-2, 2-11  
 LNS (L2TP Network Server). *See* L2TP.  
 local-address 4-52  
 LSAs  
   retransmit interval 8-11  
   type 7 8-18

## M

MAC address aging 3-9  
 MAC source address limiting 3-12  
 mac-address-learning-limit 3-3, 3-4  
 mac-entry-age-time 3-7  
 magic-number negotiation 4-101  
 management information base (MIB) for IP2000  
   statistics A-29  
 MBONE  
   Gigabit Ethernet interface 10-10  
   network-side MBONE interface 9-2  
 MBONE configuration 9-4  
 mbone-lan-interface 9-5  
 mbone-profile 9-5  
 MCAST-SERVICE 9-8  
 MD5 authentication for OSPF (RFC 2178) 8-2  
 memory modes, IP2100 1-7  
 metrics 8-5  
 mobile client connection settings 6-11  
 model numbers, IP control modules 1-3, 1-6  
 multicast  
   address range filters 9-10  
   administrative tools A-4  
   ATM QoS guidelines with unicast and  
     multicast clients 9-21  
   diagnostics A-4  
   group membership management 9-8  
   IGMP-v2 timers 9-13  
   limitations for virtual routers 7-6  
   number of clients per group 9-3  
   routed VLAN interface 9-6  
   sample configuration 9-15, 9-19  
   service profiles 9-8

transmitting streaming video 9-1, 9-2  
 video transmission 1-15  
*See also* distributed multicast mechanism  
*See also* Protocol Independent Multicast Sparse  
   Mode (PIM-SM v2)  
 multicast backbone. *See* MBONE  
 multicast server VCs 9-23  
 multicast video traffic, filtering 9-17  
 multicast-capable modules 9-26  
 multicast-server-vc 9-23  
 multiple-mcast-filter 9-11

## N

neighbors, OSPF 8-3  
 NetBIOS 4-26  
 netstat 2-4, 4-2, 4-109, 10-12  
 network alignment. *See* pools  
 network processor diagnostics A-20, A-27  
 no-bridging 3-4  
 normal areas, OSPF 8-6  
 not-so-stubby-areas (NSSAs), OSPF 8-6  
 np-default-filtering-policy 5-26  
 np-fpp-class-mem-mode 1-7  
 np-fpp-compact-timer 5-26  
 NP-PORT bandwidth allocation 1-12  
 np-update-time 5-26  
 nslookup 4-109  
 NSP VLANs 3-27  
 numbered interfaces 4-52

## O

OAM-CONFIG profile 1-8  
 Open Shortest Path First (OSPF). *See* OSPF  
 OSPF  
   ABR capability 8-2  
   adjacencies, forming 8-4  
   administrative tools 8-19  
   area border routers (ABRs) 8-6  
   ASBR calculations 8-2  
   ASE preferences, setting 8-17  
   backup designated routers (BDRs) 8-4  
   costs, configuring 8-5  
   designated router (DR) 8-4  
   MD5 authentication (RFC 2178) 8-2  
   neighbors 8-3  
   normal areas 8-6  
   not-so-stubby-areas 8-6  
   on GigE interfaces 8-10

- on WAN interfaces 8-13
- route options, configuring 8-16
- routing information 8-3
- summarized pool, importing as an ASE 8-17
- virtual interfaces, limitation 8-2
- VLSM support 8-2
- ospf-nmba-neighbor 8-15
- ospf-options 8-15

## P

- packet marking 5-3
- packet-based ECMP 4-15
- PACKET-FLOWS 5-5
- packet-type 3-41
- Password Authentication Protocol (PAP) 4-87
- performance recommendations 5-25
- per-VC queueing 1-17
- pim A-10
- PIM-GROUP-RP-MAPPING 10-5
- pim-options 10-3, 10-7
- PIM-SM
  - administrative tools A-10
- ping 2-4, 4-109
- platforms, Stinger 1-3, 1-6
- poison-reverse RIP policy 4-12
- policing algorithm 5-13
- pools
  - addresses, dynamically assigned from 4-29
  - configuring, examples of 4-26
  - global, managed by RADIPAD 4-24
  - network alignment, rules for 4-27
  - RADIPAD, specifying host 4-25
  - RADIUS 4-24, 4-26
  - route to summarized 4-29
  - slot-based address assignment 4-36
  - summarized 4-27
  - virtual routers, defined for 7-4
  - virtual routers, example of 7-7
- port blocking 3-9
- port-block-enabled 3-8
- PPP Authentication Protocol (PAP) 4-87
- PPPoA (PPP over ATM) 4-91
  - required encapsulation type 4-90
- PPPoA bridging 3-49
- PPPoA-PPPoE required encapsulation settings 4-89
- PPPoE (PPP over Ethernet) 4-95
  - on VDSL ports 4-97
  - required encapsulation type 4-89
- PPPoE Discovery packets, snooping 3-14

- PPPoE service names
  - configuration in RADIUS 4-106
  - details of negotiation 4-101
  - LLC link level 4-102
  - sample configuration 4-105
- PPPoE snooping, configuration 3-14
- PPPoE-PPPoA autodetection 4-93, 4-96
- PPPOE-SERVER 4-102
- pppoe-snooping 3-9
- prioritization
  - QoS subsystem 5-3
- priority 5-23
- priority levels, default 1-18
- profiles
  - Connection, PPP 4-101
- Protocol Independent Multicast Sparse Mode (PIM-SM v2) 10-1
- pseudo-user profiles. *See* RADIUS pseudo-user profiles
- PVCs, configuring 4-49

## Q

- QoS
  - administrative tools 5-47
  - configuration 5-4, 5-29, 5-32, 5-36
  - defaults 5-5
  - effect of fixed-rate bandwidth configuration 1-10
  - Ethernet p-bit marking 5-17
  - fragmented IP packets 5-9
  - how nonmatching packets are prioritized (the default rule) 5-11
  - interface grouping 5-8, 5-24, 5-29
  - IP and ATM 5-24
  - IP CoS subsystem optimization for tracking underlying DSL rate adaptation 5-26
  - IP ToS marking 5-17
  - layer 2 classifiers 5-6
  - monitoring 5-47
  - on VDSL ports 4-60, 4-97
  - packet classifiers 5-7
  - performance recommendations 5-25
  - policing algorithm 5-13
  - rate limiting 5-12
  - scheduling 5-12
  - tracking rate adaptation for downstream traffic shaping 5-26
- QoS classification 5-2
- QoS packet marking 5-3
- QoS prioritization 5-3
- QoS rate limiting 5-3
- QoS scheduling 5-3

qos-interface-group 5-22  
Quality of Service. *See* QoS.  
queues, depth for UDP packets 4-12  
queues, per-VC queueing 1-17

## R

RADIPAD  
    centralized pool management 4-24  
    global address pools 4-24  
radipa-hosts 4-25  
RADIUS  
    dynamic address assignment 4-30, 4-33, 4-35  
    global pools profiles 4-24  
    pools profiles 4-24  
    pools pseudo-user profiles 4-24  
    summarized pools 4-28  
RADIUS attributes 4-50, 4-57, 4-92, 4-96, 5-4, 7-7  
RADIUS profiles 4-53, 9-20  
RADIUS pseudo-user profiles  
    global-pools 4-25  
    pools 4-24  
Rapid Spanning Tree Protocol, *See* STP  
rate limiting  
    prioritization, and scheduling subsystem 5-3  
reverse-path-check 4-48  
RFC 1483 PVC  
    *See also* terminating PVCs  
RFC 2684 PVC  
    aggregation onto trunk interface 4-55  
    encapsulation methods 10-8  
    Gigabit Ethernet redundancy 2-6  
    obtaining IP address via DHCP 4-74  
    routing onto LAN 4-53  
    termination 4-51  
    *See also* terminating PVCs  
RIP 4-7  
    ignore default route in updates 4-13  
    packets, number queued 4-12  
    propagating received routes 4-12  
    triggering 4-12  
    updating changed routes only 4-12  
    virtual routers, defined for 7-4  
RLIM  
    FFW interfaces on 9-27  
routed VLANs 3-35  
routing. *See* IP routing.

## S

sar A-19  
scheduler-queue-size 5-26  
scheduling  
    increasing queue size for LIM interfaces 5-26  
    priority of the HB LIM FFW interface 9-30  
    QoS subsystem 5-3  
security  
    bridging groups isolate VLAN traffic 3-6  
    denial-of-service protection 4-9  
    filters 11-8  
    management VLANs 3-35  
    multicast group membership 9-8  
    requiring passwords for PPP sessions 4-87  
send-hostname-only 4-71  
services 1-15  
slot-based address assignment 4-36  
SNMP  
    limitations for virtual routers 7-2  
    packets, number queued 4-13  
soft IP interface configuration 4-8  
source interface local addresses 4-42  
source MAC address  
    filtering on 11-21  
    learning limit 3-12  
source-if 4-44  
split-horizon RIP policy 4-12  
stacked VLANs 3-27  
stacked-vlan 3-4  
static routes  
    assigning a cost (OSPF), example 8-19  
    OSPF, configuring 8-17  
    subnet 4-42  
    summarized pools, to 4-29  
    virtual router, defining for 7-8  
statistics  
    GigE interface A-29  
    L2TP 6-26  
    performance on LIM interfaces 4-4  
    PPPoE 6-27  
    protocol 7-5  
    traffic statistics for open UDP/TCP ports 6-26  
    VLAN A-34  
Stinger platforms that support IP control  
    modules 1-3, 1-6  
STP  
    bridge-level settings 2-19  
    configuration overview 2-18  
    limitations with this software 2-18  
    port-level settings 2-20  
    primary application with this software 2-17  
    routing implications 2-22, 2-24

- sample configurations 2-22
- stub areas, defined 8-6
- summarization. *See* pools
- summarized pool, importing as an ASE 8-17
- system IP address
  - defining 4-9
  - virtual routers, for 7-3
- system-ip-addr 4-9

## T

- TCP-established filters 11-6
- telnet 4-109
- terminating PVCs
  - ATM settings 4-49
  - IP settings 4-50
  - See also* RFC 2684 PVC 4-50
- ToS marking 5-17
- traceroute 4-109
- traffic shaping, tracking rate adaptation 5-26
- transparent-bridging 3-4
- triggering, RIP updates 4-12
- trunk terminated calls, increasing number
  - supported 1-4
- tunnel assignment IDs 6-22, 6-23
- tunnels created for clients 6-5

## U

- UDP packet queues, reducing overhead 4-12
- user connections, maximizing the number of 1-8
- User Datagram Protocol (UDP). *See* UDP
- User profiles, RADIUS. *See* RADIUS

## V

- variable-length subnet masks (VLSMs) 8-2
- VDSL
  - ETHER-QOS profile 4-60, 4-97
  - sample BIR configuration with DHCP option 82 4-82
  - sample BIR connection 4-60
  - sample PPPoE connection 4-97
- Very high speed DSL (VDSL). *See* VDSL.
- video traffic, filtering 9-17
- video, multicast 9-15
- virtual IP interfaces
  - and interface grouping 5-24

- local 4-8
- WAN, DHCP requests 4-66
- WAN, IP address lease time 4-66
- Virtual path shaping, limitation on IP2100 1-15
- virtual routers
  - address pools, for 7-4
  - configuring 7-4
  - defined 7-2
  - deleting 7-12
  - example of 7-4
  - network commands modified 7-3
  - protocol statistics 7-5
  - RIP policies 7-4
  - routed VLAN interface 3-36
  - routing table 7-5
  - static routes, defining 7-8
  - system address for 7-3

- VLAN
  - administrative tools 3-56
  - ARP table 3-48
  - backbone 3-27
  - bridge-group 3-7
  - bridging groups 3-6
  - bridging to an upstream BRAS 3-39
  - definitions 3-1
  - DHCP snooping 3-14, 3-50, 3-53
  - diagnostics A-13, A-20
  - enabling layer 2 bridging for VLAN operations 2-3
  - filters 3-4
  - filters on bridged VLANs 11-22
  - Gigabit Ethernet side of bridge circuit 3-5, 3-43, 3-44, 3-45
  - GigE redundancy 2-7
  - IGMP snooping 3-10
  - LIM interface side of bridge circuit 3-5
  - MAC address aging 3-9
  - multicast forwarding on 9-6
  - NSP VLANs 3-27
  - port blocking 3-9
  - PPPoE snooping 3-14, 3-50, 3-53
  - proxy ARP 3-46, 3-47
  - routed VLAN interface 3-35
  - stacked VLANs 3-27
  - statistics A-34
- VLAN bridging (1-to-1) 3-2
- VLAN bridging (N-to-1) 3-6
- vlan-circuit 3-4
- VLAN-ETHERNET 3-3, 3-7
- voice services 1-16

## W

- WAN OSPF interfaces

## ***Index***

### ***W***

---

- authentication 8-2
- configuring, example of 8-13
- designated router priority 8-4
- WAN virtual IP interfaces on BIR connections 4-62
- WAN virtual IP interfaces on BIR connections (DHCP) 4-66
- wan-router-interface-profile 3-8
- which 9-19