Avaya Ethernet Routing Switch 3500 Series Configuration — IP Routing and Multicast

5.0
NN47203-502
Issue 01.02
March 2013

Notice

While reasonable efforts have been made to ensure that the information in this document is complete and accurate at the time of printing, Avaya assumes no liability for any errors. Avaya reserves the right to make changes and corrections to the information in this document without the obligation to notify any person or organization of such changes.

Documentation disclaimer

“Documentation” means information published by Avaya in varying mediums which may include product information, operating instructions and performance specifications that Avaya generally makes available to users of its products. Documentation does not include marketing materials. Avaya shall not be responsible for any modifications, additions, or deletions to the original published version of documentation unless such modifications, additions, or deletions were performed by Avaya. End User agrees to indemnify and hold harmless Avaya, Avaya's agents, servants and employees against all claims, lawsuits, demands and judgments arising out of, or in connection with, subsequent modifications, additions or deletions to this documentation, to the extent made by End User.

Link disclaimer

Avaya is not responsible for the contents or reliability of any linked websites referenced within this site or documentation provided by Avaya. Avaya is not responsible for the accuracy of any information, statement or content provided on these sites and does not necessarily endorse the products, services, or information described or offered within them. Avaya does not guarantee that these links will work all the time and has no control over the availability of the linked pages.

Warranty

Avaya provides a limited warranty on its hardware and Software (“Product(s)”). Refer to your sales agreement to establish the terms of the limited warranty. In addition, Avaya’s standard warranty language, as well as information regarding support for this Product while under warranty is available to Avaya customers and other parties through the Avaya Support website: http://support.avaya.com. Please note that if you acquired the Product(s) from an authorized Avaya reseller outside of the United States and Canada, the warranty is provided to you by said Avaya reseller and not by Avaya. “Software” means computer programs in object code, provided by Avaya or an Avaya Channel Partner, whether as stand-alone products or pre-installed on hardware products, and any upgrades, updates, bug fixes, or modified versions.

Licenses

THE SOFTWARE LICENSE TERMS AVAILABLE ON THE AVAYA WEBSITE, HTTP://SUPPORT.AVAYA.COM/LICENSEINFO ARE APPLICABLE TO ANYONE WHO DOWNLOADS, USES AND/OR INSTALLS AVAYA SOFTWARE, PURCHASED FROM AVAYA INC., ANY AVAYA AFFILIATE, OR AN AUTHORIZED AVAYA RESELLER (AS APPLICABLE) UNDER A COMMERCIAL AGREEMENT WITH AVAYA OR AN AUTHORIZED AVAYA RESELLER. UNLESS OTHERWISE AGREED TO BY AVAYA IN WRITING, AVAYA DOES NOT EXTEND THIS LICENSE IF THE SOFTWARE WAS OBTAINED FROM ANYONE OTHER THAN AVAYA, AN AVAYA AFFILIATE OR AN AVAYA AUTHORIZED RESELLER; AVAYA RESERVES THE RIGHT TO TAKE LEGAL ACTION AGAINST YOU AND ANYONE ELSE USING OR SELLING THE SOFTWARE WITHOUT A LICENSE. BY INSTALLING, DOWNLOADING OR USING THE SOFTWARE, OR AUTHORIZING OTHERS TO DO SO, YOU, ON BEHALF OF YOURSELF AND THE ENTITY FOR WHOM YOU ARE INSTALLING, DOWNLOADING OR USING THE SOFTWARE (HEREINAFTER REFERRED TO INTERCHANGEABLY AS “YOU” AND “END USER”), AGREE TO THESE TERMS AND CONDITIONS AND CREATE A BINDING CONTRACT BETWEEN YOU AND AVAYA INC. OR THE APPLICABLE AVAYA AFFILIATE (“AVAYA”).

Avaya grants you a license within the scope of the license types described below, with the exception of Heritage Nortel Software, for which the scope of the license is detailed below. Where the order documentation does not expressly identify a license type, the applicable license will be a Designated System License. The applicable number of licenses and units of capacity for which the license is granted will be one (1), unless a different number of licenses or units of capacity is specified in the documentation or other materials available to you. “Designated Processor” means a single stand-alone computing device. “Server” means a Designated Processor that hosts a software application to be accessed by multiple users.

Heritage Nortel Software

“Heritage Nortel Software” means the software that was acquired by Avaya as part of its purchase of the Nortel Enterprise Solutions Business in December 2009. The Heritage Nortel Software currently available for license from Avaya is the software contained within the list of Heritage Nortel Products located at http://support.avaya.com/LicenseInfo under the link “Heritage Nortel Products”. For Heritage Nortel Software, Avaya grants Customer a license to use Heritage Nortel Software provided hereunder solely to the extent of the authorized activation or authorized usage level, solely for the purpose specified in the Documentation, and solely as embedded in, for execution on, or (in the event the applicable Documentation permits installation on non-Avaya equipment) for communication with Avaya equipment. Charges for Heritage Nortel Software may be based on extent of activation or use authorized as specified in an order or invoice.

Copyright

Except where expressly stated otherwise, no use should be made of materials on this site, the Documentation, Software, or hardware provided by Avaya. All content on this site, the documentation and the Product provided by Avaya including the selection, arrangement and design of the content is owned either by Avaya or its licensors and is protected by copyright and other intellectual property laws including the sui generis rights relating to the protection of databases. You may not modify, copy, reproduce, republish, upload, post, transmit or distribute in any way any content, in whole or in part, including any code and software unless expressly authorized by Avaya. Unauthorized reproduction, transmission, dissemination, storage, and or use without the express written consent of Avaya can be a criminal, as well as a civil offense under the applicable law.

Third Party Components

“Third Party Components” mean certain software programs or portions thereof included in the Software that may contain software (including open source software) distributed under third party agreements (“Third Party Components”), which contain terms regarding the rights to use certain portions of the Software (“Third Party Terms”). Information regarding distributed Linux OS source code (for those Products that have distributed Linux OS source code) and identifying the copyright holders of the Third Party Components and the Third Party Terms that apply is available in the Documentation or on Avaya’s website at: http://support.avaya.com/Copyright. You agree to the Third Party Terms for any such Third Party Components.

Preventing Toll Fraud

“Toll Fraud” is the unauthorized use of your telecommunications system by an unauthorized party (for example, a person who is not a corporate employee, agent, subcontractor, or is not working on your company's behalf). Be aware that there can be a risk of Toll Fraud associated with your system and that, if Toll Fraud occurs, it can result in substantial additional charges for your telecommunications services.

Avaya Toll Fraud intervention

If you suspect that you are being victimized by Toll Fraud and you need technical assistance or support, call Technical Service Center Toll Fraud Intervention Hotline at +1-800-643-2353 for the United States and Canada. For additional support telephone numbers, see the Avaya Support website: http://support.avaya.com. Suspected security

vulnerabilities with Avaya products should be reported to Avaya by sending mail to: securityalerts@avaya.com.

Trademarks

The trademarks, logos and service marks (“Marks”) displayed in this site, the Documentation and Product(s) provided by Avaya are the registered or unregistered Marks of Avaya, its affiliates, or other third parties. Users are not permitted to use such Marks without prior written consent from Avaya or such third party which may own the Mark. Nothing contained in this site, the Documentation and Product(s) should be construed as granting, by implication, estoppel, or otherwise, any license or right in and to the Marks without the express written permission of Avaya or the applicable third party.

Avaya is a registered trademark of Avaya Inc.

All non-Avaya trademarks are the property of their respective owners, and “Linux” is a registered trademark of Linus Torvalds.

Downloading Documentation

For the most current versions of Documentation, see the Avaya Support website: http://support.avaya.com.

Contact Avaya Support

See the Avaya Support website: http://support.avaya.com for product notices and articles, or to report a problem with your Avaya product. For a list of support telephone numbers and contact addresses, go to the Avaya Support website: http://support.avaya.com, scroll to the bottom of the page, and select Contact Avaya Support.

Contents

Chapter 1: Purpose of this document

This document provides procedures and conceptual information to configure IP routing features on the ERS 3500 Series, including static routes, Proxy ARP, DHCP Relay, and UDP forwarding. It also provides procedures and conceptual information to manage multicast traffic using IGMP snooping.

Chapter 2: New in this release

This is a new document for Avaya Ethernet Routing Switch 3500 Series Release 5.0.

The Avaya ERS 3500 Series is new and supports the following hardware and software features:

ERS 3500 hardware

The following table lists and describes the supported hardware for ERS 3500 Series 5.0. Question marks (?) in the table signify power cord types; substitute the following regional variants:

- A — no power cord
- B — EU power cord
- C — UK / Ireland power cord
- D — Japan power cord
- E — North American power cord
- F — Australia / New Zealand / China power cord

Note:

All switches support autopolarity.

Table 1: Hardware

Hardware	Description
Switch models	
AL3500?01–E6	3526T — 24 10/100BaseT ports supporting autosensing and autonegotiation, in a non-PoE , plus two 10/100/1000 or Small Form Pluggable (SFP) front combination ports, plus two SFP rear ports. Fanless.
AL3500?11–E6	3526T-PWR+ — 24 10/100BaseT PoE+ ports (802.3af/at), plus two 10/100/1000 or Small Form Pluggable (SFP) front combination ports, plus two SFP rear ports.
AL3500?04–E6	3510GT — 8 10/100/1000BaseT ports, plus two SFP ports (ports 9 and 10). Standalone and fanless.
AL3500?14–E6	3510GT-PWR+ — 8 10/100/1000BaseT PoE+ ports (802.3af/at), plus two SFP ports (ports 9 and 10). Standalone. Fanless operation in Low Power mode @ 60W max PoE budget, or normal fan

Hardware	Description
	operation in High Power mode @ 170W max PoE budget.
AL3500?05–E6	3524GT — 24 10/100/1000BaseT ports, four SFP ports shared with ports 21–24, plus two SFP rear ports.
AL3500?15–E6	3524GT-PWR+ — 24 10/100/1000BaseT PoE+ ports (802.3af/at), four SFP ports shared with ports 21–24, plus two SFP rear ports.
Rack Mount Kits	
AL3511001–E6	Spare Rack Mount Kit — this kit can be used as a replacement rack mount kit for ERS 3524GT, ERS 3524GT-PWR+, ERS 3526T or ERS 3526T-PWR + switches.
AL3511002–E6	3510–Pair Rack Mount Kit — this kit is used to connect two ERS 3510GT or ERS 3510GT-PWR+ switches together side by side and mount them in a 19 inch rack.
AL3511003–E6	3510–Single Rack Mount Kit — this kit is used to mount a single ERS 3510GT or ERS 3510GT-PWR+ switch in a standard 19 inch rack.

ERS 3500 software features

The following software features are supported on the ERS 3500 Series Release 5.0:

- BootP or Default IP
- RADIUS password fallback
- Downloading agent & diags without reset
- Username Password enhancement
- Autosave configuration enhancements
- Ping enhancement
- Writemem and save config command
- Configurable SNMP trap port (only SNMP v1 & v2)
- SNTP & SNTP timezone enhancement
- Shutdown, reload enhancement
- Factory-default command
- Show MAC address enhancement
- Show Port enhancement
- Show Running Config (verbose, non-verbose, module) enhancement

- VLAN Tagging enhancement
- 802.1AB (LLDP) Standards Based Auto Topology
- 802.1w&s — rapid and multiple spanning trees
- 802.3ad- Link Aggregation Control Protocol (LACP)
- 802.3af — Power over Ethernet (PoE)
- 802.3at — Power over Ethernet plus (PoE+)
- COS/DSCP — allows mapping the DSCP value (carried by IP frames) to 802.1p priority value
- Rate Limiting
- Remote logging — ability to log on remote servers
- Web Quick Start
- WEB HTTP download of ASCII — allows downloading of ASCII configuration files through HTTP
- HTTP web-based management
- HTTPS/SSL secure web management
- HTTP port change
- CLI Quick Start script
- Auto save Disable
- Telnet (up to four sessions)
- Telnet out — ability to open Telnet sessions from the box
- Domain Name Service (DNS) capability
- 256 port-based VLANs with IVL — VLAN 1 is the default management VLAN
- 802.1Q tagging
- 802.1p traffic class support / remarking
- Advanced QoS (traffic classification, filtering, mark/remarking, metering, shaping)
- Avaya Automatic-QoS
- Single 802.1d Spanning Tree Protocol (STP) on all ports
- Spanning Tree port mode
- Spanning Tree 802.1d compliance mode
- Port mirroring (1–1)
- Multi-Link Trunking (MLT) with up to six trunks and four links per trunk
- MLT enable/disable whole trunk
- IGMP Multicast no flood command enhancements
- IGMPv1/v2 snooping / proxy
- IGMPv3 Snooping/proxy
- MAC address based security with autolearn (BaySecure)

- Sticky MAC
- RADIUS-based security
- TACACS+
- Local password protection
- SNMPv3 security
- SNMP-based network management
- SNMP MIB web page in EDM
- SNMP Trap list web page in EDM
- Extended IP Manager (IPv4 & IPv6)
- IPv6 Management
- IPv6 VLANs (protocol based)
- No Banner & CTRL-Y Skip
- Local console via serial interface
- 802.3x (Flow Control — Gig ports only)
- BootP/TFTP for downloading software and config file
- RMON (RFC1757): per port Statistics, History, Alarm and Events
- ASCII file configuration
- Syslog
- Dual Syslog servers
- ASCII Config Generator (ACG)
- 802.1X EAP (SHSA, MHMA, MHSA, Guest VLAN, Non-EAP & RADIUS MAC)
- 802.1X Enhancement: Dynamic VLAN assignment for NEAP & MHMA
- 802.1X Enhancement: Unicast request, Non-EAP IP Phone support
- 802.1X RFC3576 RADIUS auth extensions - CoA
- 802.1X RFC2866/2869 RADIUS interim accounting updates
- 802.1X NEAP with VLAN names
- 802.1X NEAP last assigned VLAN
- 802.1X NEAP fail-open VLAN
- 802.1X NEAP re-authentication timer
- 802.1X NEAP and Guest VLAN on same port
- RADIUS EAP / NEAP to different servers
- RADIUS Server reachability
- DA Filtering
- Port Naming

- CANA
- SSHv2
- SSH enhancement to support RSA
- Secure FTP (SFTP)
- Auto Detection And Configuration (ADAC) with 802.1AB interaction
- 802.1AB MED (Cisco IP Phones)
- 802.1AB Location TLV
- 802.1AB and ADAC interoperability
- 802.1AB Integration features
- 802.1AB Customization features
- Identify Units (Blink LEDs)
- Cumulative system uptime (hidden command)
- Virtual LACP
- Static Routing with default route
- IP Local and Non-Local static routing
- BootP/DHCP Relay
- Proxy ARP
- UDP forwarding
- DHCP Snooping
- DHCP Client
- DHCP Option 82
- Dynamic ARP Inspection
- IP Source Guard
- BDPU Filtering
- MAC flush
- Software Exception Log
- CPU & Memory Utilization
- Configure Asset ID
- Show environmental
- Show software status

Chapter 3: Introduction

This document provides procedures and conceptual information to configure IP routing features on the Avaya Ethernet Routing Switch 3500 Series, including static routes, Proxy ARP, DHCP Relay, and UDP forwarding. It also provides procedures and conceptual information to manage multicast traffic using IGMP snooping.

ACLI command modes

Avaya command line interface (ACLI) provides the following configuration modes:

- User EXEC
- Privileged EXEC
- Global Configuration
- Interface Configuration Mode

Mode access is determined by access permission levels and password protection.

If no password is set, you can enter ACLI in User EXEC mode and use the enable command to move to the next level (Privileged EXEC mode). However, if you have read-only access, you cannot progress beyond User EXEC mode, the default mode. If you have read-write access you can progress from the default mode through all of the available modes.

With sufficient permission, you can use the rules in the following table to move between the command modes.

Command mode and sample prompt	Entrance commands	Exit commands
User EXEC `3524GT>`	No entrance command, default mode.	Type `exit`or `logout`
Privileged EXEC `3524GT#`	From User EXEC mode, type: `enable`	Type `exit` or `logout`
Global Configuration `3524GT(config)#`	From Privileged EXEC mode, type: `configure`	To return to Privileged EXEC mode, type: `end` or `exit` To exit ACLI completely, type: `logout`

Command mode and sample prompt	Entrance commands	Exit commands

Interface Configuration
`3524GT(config-if)#`
	From Global Configuration mode:
	To configure a port, type:
	`interface fastethernet <port number>`
	To configure a VLAN, type:
	`interface vlan <vlan number>`
		To return to Global Configuration mode, type:
		`exit`
		To return to Privileged EXEC mode, type: `end`
		To exit ACLI completely, type:
		`logout`

For more information about the ACLI configuration modes, see *Avaya Ethernet Routing Switch 3500 Series Fundamentals* (NN47203-102).

Chapter 4: IP routing fundamentals

This chapter provides an introduction to IP routing and related features used in the Avaya Ethernet Routing Switch 3500 Series.

IP addressing overview

An IP version 4 (IPv4) address consists of 32 bits expressed in a dotted-decimal format (XXX.XXX.XXX.XXX). The IPv4 address space is divided into classes, with classes A, B, and C reserved for unicast addresses, and accounting for 87.5 percent of the 32-bit IP address space. Class D is reserved for multicast addressing. The following table lists the breakdown of the IP address space by address range and mask.

Class	Address Range	Mask	Number of Networks	Nodes per Network
A	1.0.0.0 - 127.0.0.0	255.0.0.0	127	16 777 214
B	128.0.0.0 - 191.255.0.0	255.255.0.0	16 384	65 534
C	192.0.0.0 - 223.255.255.0	255.255.255.0	2 097 152	255
D	224.0.0.0 - 239.255.255.254			
E	240.0.0.0 - 240.255.255.255			

Note:

Class D addresses are primarily reserved for multicast operations, although the addresses 224.0.0.5 and 224.0.0.6 are used by OSPF and 224.0.0.9 is used by RIP

Note:

Although technically part of Class A addressing, network 127 is reserved for loopback.

Note:

Class E addresses are reserved for research purposes.

To express an IP address in dotted-decimal notation, each octet of the IP address is converted to a decimal number and separated by decimal points. For example, the 32-bit IP address

10000000 00100000 00001010 10100111 is expressed in dotted-decimal notation as 128.32.10.167.

Each IP address class, when expressed in binary notation, has a different boundary point between the network and host portions of the address, as shown in the following figure. The network portion is a network number field from 8 through 24 bits. The remaining 8 through 24 bits identify a specific host on the network.

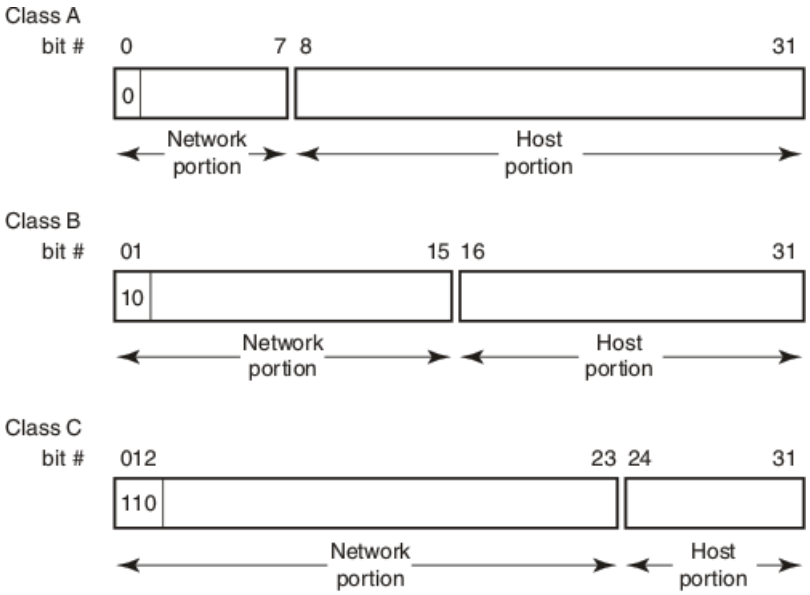

Figure 1: Network and host boundaries in IP address classes

Subnet addressing

Subnetworks (or subnets) are an extension of the IP addressing scheme. With subnets, organizations can use one IP address range for multiple networks. Subnets are two or more physical networks that share a common network-identification field (the network portion of the 32-bit IP address).

A subnet address is created by increasing the network portion to include a subnet address, thus decreasing the host portion of the IP address. For example, in the address 128.32.10.0, the network portion is 128.32, while the subnet is found in the first octet of the host portion (10). A subnet mask is applied to the IP address and identifies the network and host portions of the address.

The following table illustrates how subnet masks used with Class B and Class C addresses can create differing numbers of subnets and hosts. This example shows the use of the zero subnet, which is permitted on a Avaya Ethernet Routing Switch 3500 Series.

Number of bits	Subnet Mask	Number of Subnets (Recommended)	Number of Hosts per Subnet
Class B			
2	255.255.192.0	2	16 382
3	255.255.224.0	6	8190
4	255.255.240.0	14	4094
5	255.255.248.0	30	2046
6	255.255.252.0	62	1022
7	255.255.254.0	126	510
8	255.255.255.0	254	254
9	255.255.255.128	510	126
10	255.255.255.192	1022	62
11	255.255.255.224	2046	30
12	255.255.255.240	4094	14
13	255.255.255.248	8190	6
14	255.255.255.252	16 382	2
Class C			
1	255.255.255.	0	126
2	255.255.255.192	2	62
3	255.255.255.224	6	30
4	255.255.255.240	14	14
5	255.255.255.248	30	6
6	255.255.255.252	62	2

Variable-length subnet masking (VLSM) is the ability to divide an intranet into pieces that match network requirements. Routing is based on the longest subnet mask or network that matches.

IP routing

To configure IP routing on the Avaya Ethernet Routing Switch 3500 Series, you must create virtual router interfaces by assigning an IP address to a virtual local area network (VLAN). The following sections provide more details about IP routing functionality.

For a more detailed description about VLANs and their use, see *Avaya Ethernet Routing Switch 3500 Series, Configuration - Layer 2* (NN47215–500).

IP routing using VLANs

The Avaya Ethernet Routing Switch 3500 Series, supports wire-speed IP routing between VLANs. To create a virtual router interface for a specified VLAN, you must associate an IP address with the VLAN.

The virtual router interface is not associated with any specific port. The VLAN IP address can be reached through any of the ports in the VLAN. The assigned IP address also serves as the gateway through which packets are routed out of that VLAN. Routed traffic can be forwarded to another VLAN within the switch or stack.

When the Avaya Ethernet Routing Switch 3500 Series, is routing IP traffic between different VLANs, the switch is considered to be running in Layer 3 mode; otherwise, the switch runs in Layer 2 mode. When you assign an IP address to a Layer 2 VLAN, the VLAN becomes a routable Layer 3 VLAN.

You can assign a single and unique IP address to each VLAN. You can configure the global status of IP routing to be enabled or disabled on the Avaya Ethernet Routing Switch 3500 Series,. By default, IP routing is disabled.

In this release, the Avaya Ethernet Routing Switch 3500 Series supports local routes and static routes (local and non-local static routes). With local routing, the switch automatically creates routes to each of the local Layer 3 VLAN interfaces. With static routing, you must manually enter the routes to the destination IP addresses.

Local routes

With routing globally enabled, if you assign an IP address to a VLAN, IP routing is enabled for that VLAN. In addition, for each IP address assigned to a VLAN interface, the Ethernet Routing Switch adds a directly connected or local route to its routing table based on the IP address/mask assigned.

Local routing example

The following figure shows how the Ethernet Routing Switch can route between Layer 3 VLANs. In this example, the Ethernet Routing Switch has two VLANs configured. IP Routing is enabled globally on the switch and on the VLANs, each of which has an assigned IP address.

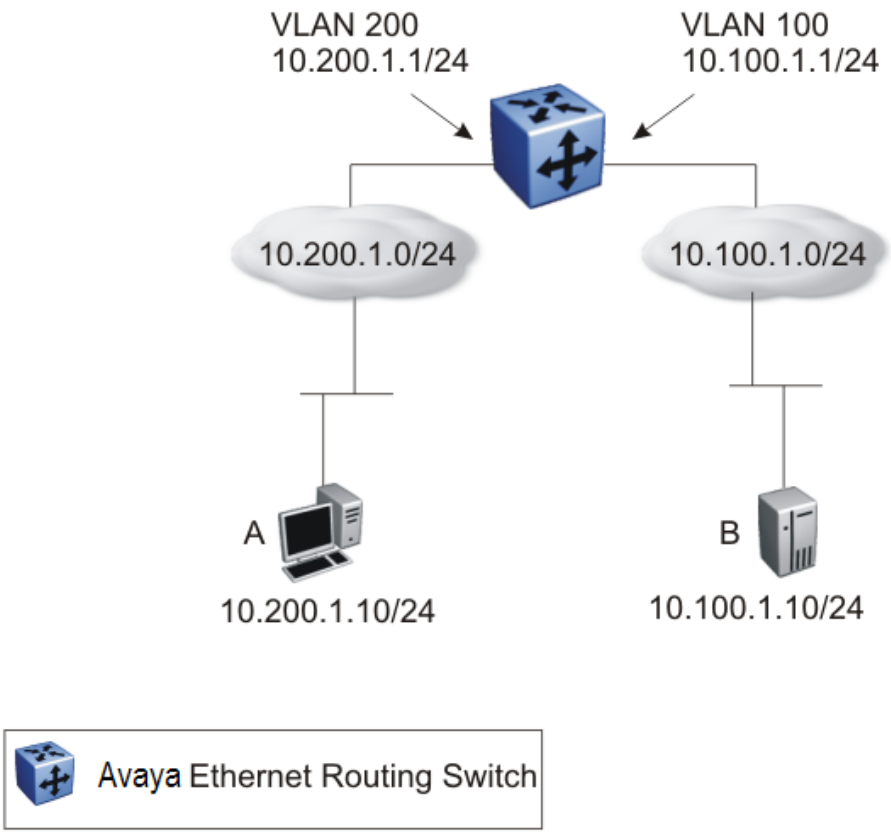

Figure 2: Local routes example

IP address 10.100.1.1/24 is assigned to VLAN 100, and IP address 10.200.1.1/24 is assigned to VLAN 200. As IP Routing is enabled, two local routes become active on the Avaya Ethernet Routing Switch as described in the following table.

	Network	Net-mask	Next-hop	Type
1	10.100.1.0	255.255.255.0	10.100.1.1	LOCAL
2	10.200.1.0	255.255.255.0	10.200.1.1	LOCAL

At this stage, both hosts A (10.200.1.10) and B (10.100.1.10) are reachable from the Ethernet Routing Switch. However, to achieve Layer 3 connectivity between A and B, additional configuration is required. Host A must know how to reach network 10.100.1.0/24, and host B must know how to reach network 10.200.1.0/24.

On host A, you must configure a route to network 10.100.1.0/24 through 10.200.1.1, or configure 10.200.1.1 as the default gateway for the host.

On host B, you must configure a route to network 10.200.1.0/24 through 10.100.1.1, or configure 10.100.1.1 as the default gateway for the host.

With these routes configured, the Ethernet Routing Switch can perform inter-VLAN routing, and packets can flow between hosts A and B.

Static routes

After you create routable VLANs though IP address assignment, you can create static routes. With static routes, you can manually create specific routes to a destination IP address. In this release, the Ethernet Routing Switch supports both local and non-local static routes.

Static routes are not easily scalable. Thus, in a large or growing network, this type of route management may not be optimal.

Static routing example

The following figure shows an example of static routing on the Ethernet Routing Switch.

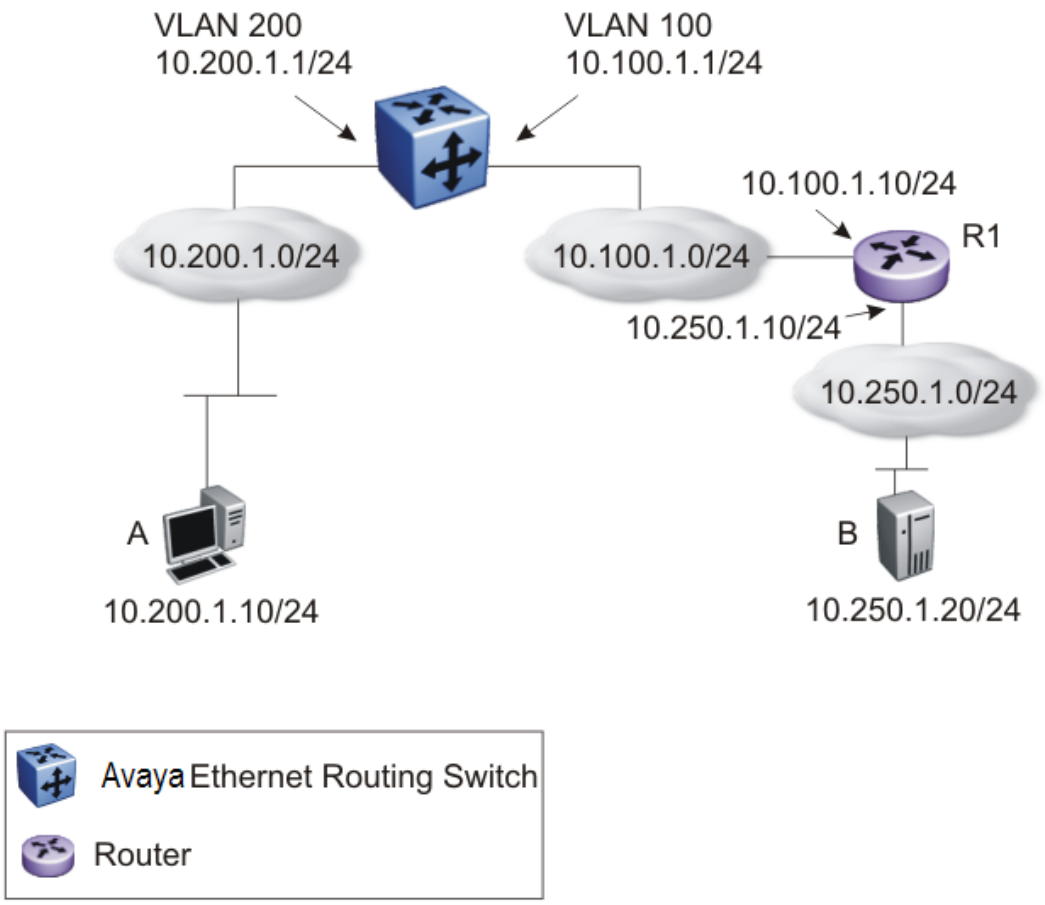

Figure 3: Static routes

In this example, two Layer 3 devices are used to create a physical link between hosts A and B. This network contains an Ethernet Routing Switch and another Layer 3 router, R1.

In this setup, the local route configuration from Local routing example on page 22 still applies. However, in this case, network 10.100.1.0/24 stands in between networks 10.200.1.0/24 and 10.250.1.0/24. To achieve end-to-end connectivity, router R1 must know how to reach network 10.200.1.0/24, and the Ethernet Routing Switch must know how to reach network 10.250.1.0/24. On the Ethernet Routing Switch, you can accomplish this using static routing. With static routing, you can configure a route to network 10.250.1.0/24 through 10.100.1.10. In this case, the following routes are active on the Ethernet Routing Switch.

	Network	Net-mask	Next-hop	Type
1	10.100.1.0	255.255.255.0	10.100.1.1	LOCAL
2	10.200.1.0	255.255.255.0	10.200.1.1	LOCAL
3	10.250.1.0	255.255.255.0	10.100.1.10	STATIC

To obtain Layer 3 connectivity between the hosts, additional routes are required. Host A requires a route to 10.250.1.0/24 using 10.200.1.1 as the next hop, or with 10.200.1.1 as the

default gateway. Host B requires a route to 10.200.1.0/24 using 10.250.1.10 as the next hop, or with 10.250.1.10 as the default gateway.

The configuration for router R1 to reach network 10.200.1.0/24 is dependent on the type of router used.

Layer 3 Non-Local Static Routes (IP NLSR)

After you create routable VLANs through IP address assignment, you can create static routes.

You can manually create specific routes to destination IP addresses with static routes.

Local static routes have a next-hop that is on a directly-connected network.

Non-local routes (NLSR) have a next-hop that is not on a directly-connected network.

When you implement NLSR on the switch, if the corresponding next-hop IP address can be reached through any active route on the switch, a static route becomes active in the routing table.

The switch elects a supported route as the most specific route through which the next-hop IP address can be reached. Then the switch links the NLSR route to an active supported route. The NLSR becomes inactive if the supported route becomes inactive and no alternative supported route can be calculated.

The supported route can be a static route or dynamic route (on switches that support dynamic routing), but it cannot be the default route (network 0.0.0.0 netmask 0.0.0.0) because, if NLSR reachability is allowed through the default route, then any route could change to active as NLSR reachable through the default route.

Advantages of IP NLSR:

- Where there are multiple paths to a network you can reduce the number of static routes by using only one route with a remote gateway
- Where the next-hop IP address cannot be reached directly from the switch, the system can use any host IP address that exists on the path to the destination network to configure an active and functional route, as long as the host can be reached through another active route on the switch
- You do not need to modify the NLSR route if an administrator changes the next-hop IP address
- If the supported route is an ECMP route, and one of the next-hops becomes unreachable, the NLSR route remains active as long as the support route is active through at least one of the next-hops
- If the supported route is an ECMP route, internally, the NLSR route uses the first of the ECMP route next-hops as the NLSR next-hop

Limitations of IP NLSR:

- Because static routes are not easily scalable, in a large or growing network this type of route management may not be the best option
- Because static routes cannot determine path failure, a router can still attempt to use a failed path

Default routes

Default routes specify a route to all networks for which there are no explicit routes in the Forwarding Information Base or the routing table. This static default route is a route to the network address 0.0.0.0 as defined by the Institute of Electrical and Electronics Engineers (IEEE) Request for Comment (RFC) 1812 standard.

The Ethernet Routing Switch uses the default route 0.0.0.0/0.0.0.0 for all Layer 3 traffic that does not match a specific route. This traffic is forwarded to the next-hop IP address specified in the default route.

Route scaling

The Avaya Ethernet Routing Switch 3500 Series supports a maximum of 32 local routes and up to 32 static routes, including the default route (Destination = 0.0.0.0, Mask = 0.0.0.0).

Management VLAN

With IP routing enabled on the switch or stack, you can use any of the virtual router IP addresses for device management over IP. Any routable Layer 3 VLAN can carry the management traffic for the switch, including Telnet, Simple Network Management Protocol (SNMP), BootP, and Trivial File Transfer Protocol (TFTP). Without routing enabled, the management VLAN is reachable only through the switch or stack IP address, and only through ports that are members of the management VLAN. The management VLAN always exists on the switch and cannot be removed.

When routing is enabled on the Avaya Ethernet Routing Switch 3500 Series, switches, the management VLAN behaves similar to other routable VLANs. The IP address is reachable through any virtual router interface, as long as a route is available.

Management route

On the Ethernet Routing Switch, you can configure a management route from the Management VLAN to a particular subnet. The management route is a static route that allows incoming management connections from the remote network to the management VLAN.

The management route transports traffic between the specified destination network and the Management VLAN only. It does not carry inter-VLAN routed traffic from the other Layer 3 VLANs to the destination network. This provides a management path to the router that is inaccessible from the other Layer 3 VLANs. While you can access the management VLAN from all static routes, other static routes cannot route traffic to the management route.

To allow connectivity through a management route, you must enable IP routing globally and on the management VLAN interface.

The following figure shows an example of a management route allowing access to the management VLAN interface.

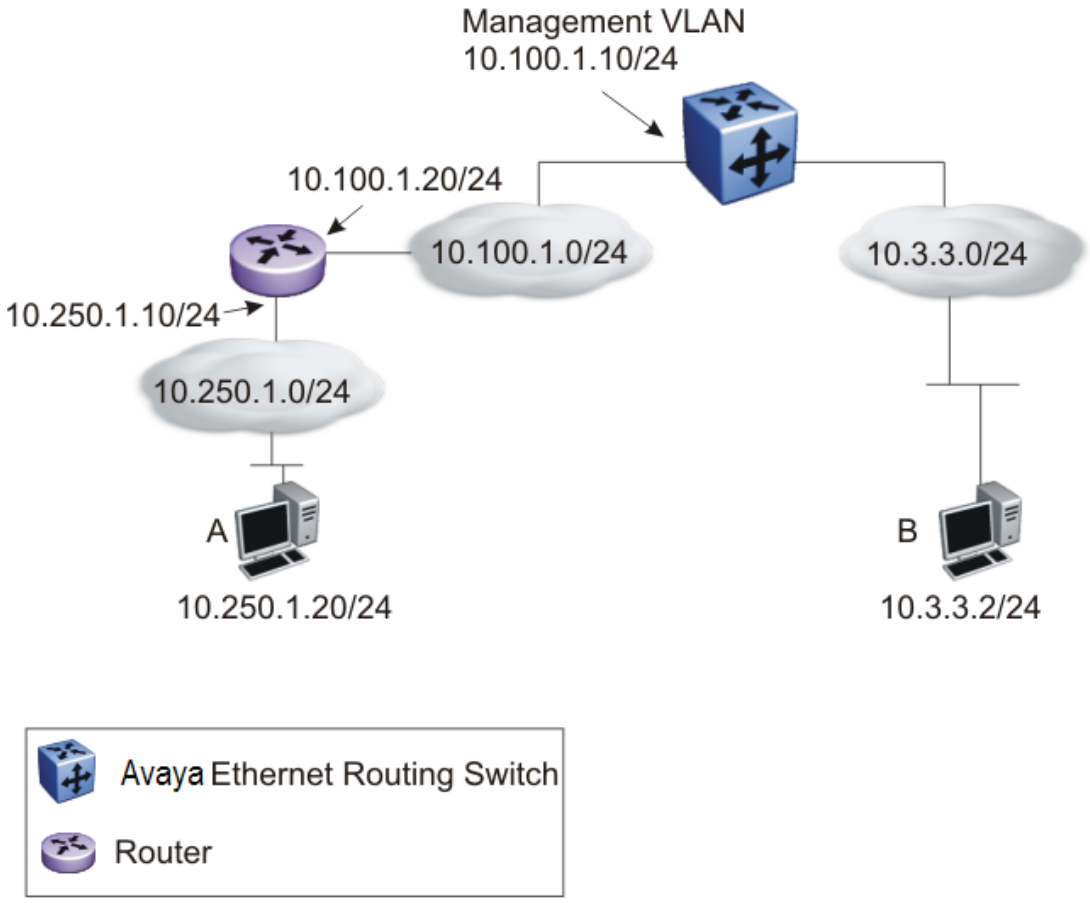

Figure 4: Management route

As network 10.250.1.0/24 is not directly connected to the Ethernet Routing Switch, to achieve connectivity from host 10.250.1.20 to the management VLAN, the Ethernet Routing Switch must know how to reach network 10.250.1.0/24. On the Ethernet Routing Switch, you can configure a management route to network 10.250.1.0/24 through 10.100.1.20. In this case, the following management route is active on the Ethernet Routing Switch.

	Network	Net-mask	Next-hop	Type
1	10.250.1.0	255.255.255.0	10.100.1.20	MANAGEMENT

With this configured route, host A at 10.250.1.20 can perform management operations on the Ethernet Routing Switch. To do so, Host A also requires a route to 10.100.1.0/24 using 10.250.1.10 as the next hop, or with 10.250.1.10 as the default gateway.

If a Layer 3 VLAN is also configured for network 10.3.3.0/24, this provides a local route that host B at 10.3.3.2 can use to access the switch. However, host B cannot communicate with host A, as the route to network 10.250.1.0/24 is a management route only. To provide connectivity between the two hosts, you must configure a static route to 10.250.1.0/24.

Related routing features

The following sections describe features that are related to and dependent on the IP routing functionality.

BootP DHCP relay

Dynamic Host Configuration Protocol (DHCP) is a mechanism to assign network IP addresses on a dynamic basis to clients who request an address. DHCP is an extension of the Bootstrap protocol (BootP). BootP/DHCP clients (workstations) generally use User Datagram Protocol (UDP) broadcasts to determine their IP addresses and configuration information. If such a host is on a VLAN that does not include a DHCP server, the UDP broadcasts are by default not forwarded to servers located on different VLANs.

The Avaya Ethernet Routing Switch 3500 Series, can resolve this issue using DHCP relay, which forwards the DHCP broadcasts to the IP address of the DHCP server. Network managers prefer to configure a small number of DHCP servers in a central location to lower administrative overhead. Routers must support DHCP relay so that hosts can access configuration information from servers several router hops away.

With DHCP relay enabled, the switch can relay client requests to DHCP servers on different Layer 3 VLANs or in remote networks. It also relays server replies back to the clients.

To relay DHCP messages, you must create two Layer 3 VLANs: one connected to the client and the other providing a path to the DHCP server. You can enable DHCP relay on a per-VLAN basis.

The following figure shows a DHCP relay example, with an end station connected to subnet 1, corresponding to VLAN 1. The Avaya Ethernet Routing Switch 3500 Series, connects two subnets by means of the virtual routing function. When the end station generates a DHCP request as a limited UDP broadcast to the IP address of all 1s (that is, 255.255.255.255), with

the DHCP relay function enabled, the Ethernet Routing Switch forwards the DHCP request to the host address of the DHCP server on VLAN 2.

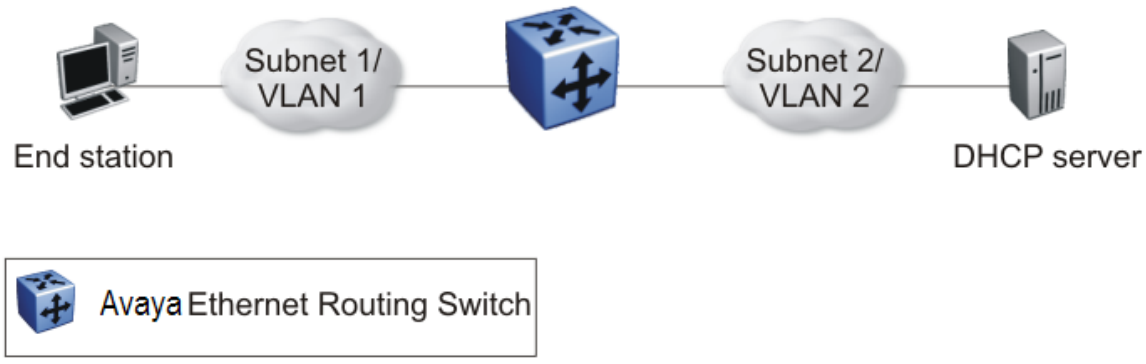

Figure 5: DHCP relay operation

Forwarding DHCP packets

In the following figure, the DHCP relay agent address is 10.10.1.254. To configure the Avaya Ethernet Routing Switch 3500 Series,to forward DHCP packets from the end station to the server, use 10.10.2.1 as the server address.

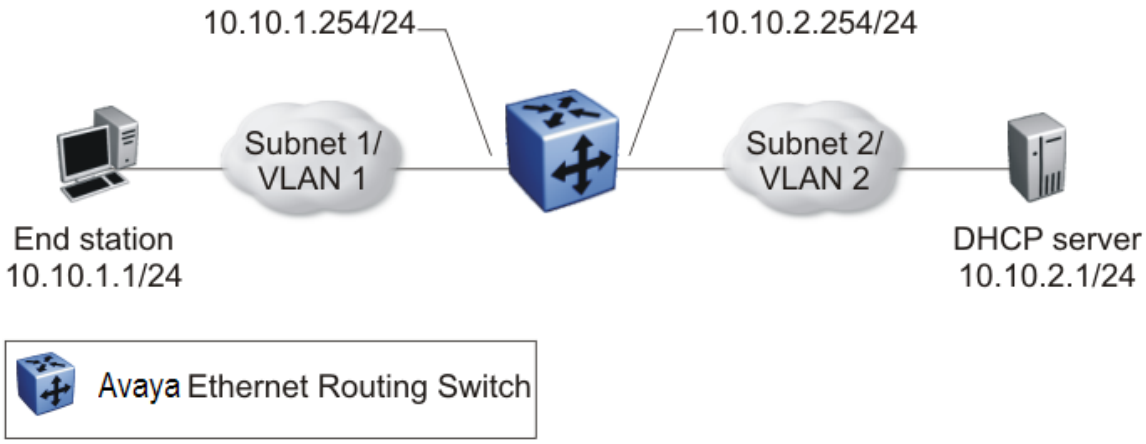

Figure 6: Forwarding DHCP packets

All BootP and DHCP broadcast packets that appear on the VLAN 1 router interface (10.10.1.254) are then forwarded to the DHCP server. In this case, the DHCP packets are forwarded as unicast to the DHCP server IP address.

Differences between DHCP and BootP

With DHCP relay, the Avaya Ethernet Routing Switch 3500 Series, supports the relay of DHCP and the Bootstrap protocol (BootP). The following differences between DHCP and BootP are specified in RFC 2131:

- BootP enables the retrieval of an American Standard Code for Information Interchange (ASCII) configuration file name and configuration server address.
- A properly configured BootP server enables the switch to automatically learn its assigned IP address, subnet mask, and the IP address of the default router (default gateway).
- DHCP defines mechanisms through which clients can be assigned a network address for a finite lease (allowing for reuse of IP addresses).
- DHCP provides the mechanism for clients to acquire all of the IP configuration parameters they need to operate.

DHCP uses the BootP message format defined in RFC 951. The remainder of the options field consists of a list of tagged parameters that are called options(RFC 2131).

DHCP option 82 support

DHCP option 82 support is an extension of Dynamic Host Configuration Protocol (RFC3046 and RFC3993) that enables the switch to send information about DHCP clients to the authenticating DHCP server. When you enable option 82, in either Layer 2 or Layer 3 mode, the switch inserts additional port-based identification information into the DHCP packets traversing the switch enroute to the DHCP server. The DHCP server stores this additional identification information within the IP allocation record to assist in tracking of end device locations; for example, to provide location-based information for emergency services applications.

When a VLAN is operating in Layer 2 mode, DHCP Snooping must be enabled for DHCP Option 82 to function, both globally and on each client VLAN. For more information about DHCP Snooping, see *Avaya Ethernet Routing Switch 3500 Series Configuration, Security* (NN47203-504).

When a VLAN is operating in Layer 3 (IP Routing) mode, the DHCP Option 82 function requires that DHCP Relay is appropriately configured. To use DHCP Option 82 with DHCP relay, you must enable DHCP relay globally on the switch and client VLANs. And you must configure at least one forward path.

If you configure two DHCP Servers (one in the same VLAN with the DHCP Client and one in another VLAN) and you enable both DHCP Snoooping Option 82 and DHCP Relay Option 82, the system adds the option for both servers.

DHCP Relay Packet Size

In accordance with RFC3046, you can specify the maximum frame size the DHCP relay agent can forward to the DHCP server. While the switch implementation permits configuration of the

maximum DHCP packet size up to 1536 bytes, the default maximum size is 576 bytes. If the DHCP frame received is larger that the configured frame size, the switch does not relay the packet. If the DHCP packet exceeds the maximum configured size, the DHCP Option 82 information is not appended to the message.

UDP broadcast forwarding

By default, User Datagram Protocol (UDP) broadcast frames received on one VLAN are not routed to another VLAN. To allow UDP broadcasts to reach a remote server, the Ethernet Routing Switch supports UDP broadcast forwarding, which forwards the broadcasts to the server through a Layer 3 VLAN interface.

UDP broadcast forwarding is a general mechanism for selectively forwarding limited UDP broadcasts received on an IP interface to a configured IP address. The packet is sent as a unicast packet to the server.

When a UDP broadcast is received on a router interface, it must meet the following criteria to be considered for forwarding:

- It must be a MAC-level broadcast.
- It must be an IP-limited broadcast.
- It must be for a configured UDP protocol.
- It must have a time-to-live (TTL) value of at least 2.

For each ingress interface and protocol, the UDP broadcast packets are forwarded only to a unicast host address (for example, to the unicast IP address of the server).

When the UDP forwarding feature is enabled, a filter is installed that compares the UDP destination port of all packets against all the configured UDP forwarding entries. If a match occurs, the destination IP of the incoming packet is checked for consistency with the userconfigured broadcast mask value for this source VLAN. If these conditions are met, the TTL field from the incoming packet is overwritten with the user-configured TTL value, the destination IP of the packet is overwritten with the configured destination IP, and the packet is routed to the destination as a unicast frame.

UDP forwarding example

The following figure shows an example of UDP broadcast forwarding. In this case, if host A (10.200.1.10) needs a certain service (for example, a custom application that listens on UDP port 12345), it transmits a UDP broadcast frame. By default, the Ethernet Routing Switch does not forward this frame to VLAN 100, and because server B (10.100.1.10) is not on VLAN 200, the host cannot access that service.

With UDP broadcast forwarding enabled, the host can access the service. In this case, you must list port 12345 as a valid forwarding port, and specify VLAN 200 as the source VLAN.

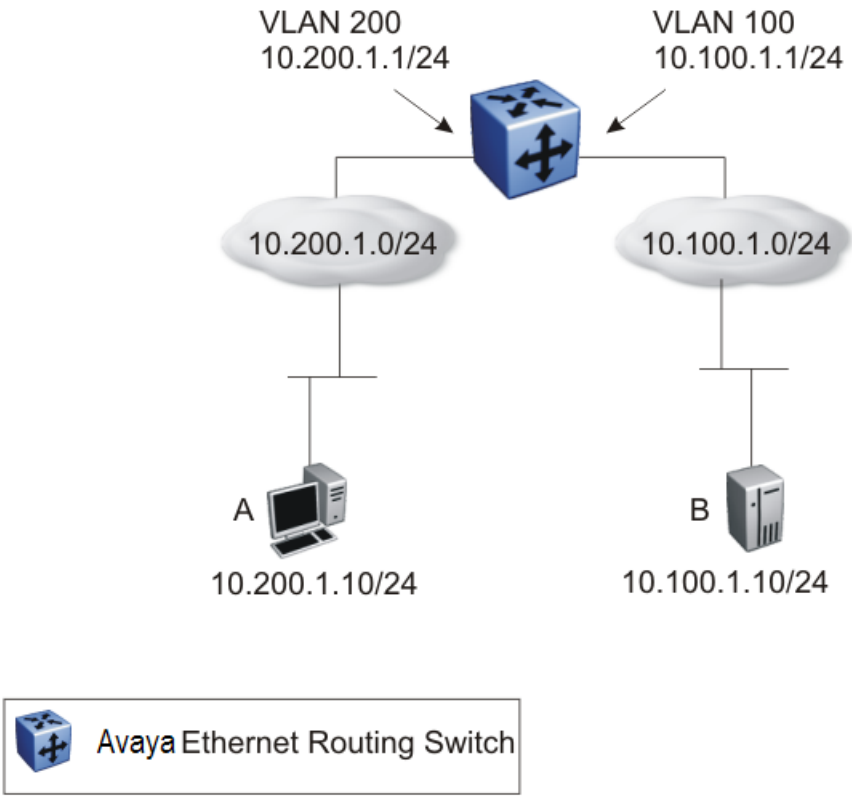

Figure 7: UDP forwarding example

When the switch receives an incoming packet on VLAN 200 that matches the configured UDP destination port (12345), and the destination IP is consistent with the broadcast mask value for the VLAN, then the switch applies the new destination IP (here, 10.100.1.10) to the packet and routes it to the destination as a unicast frame.

Directed broadcasts

With the directed broadcasts feature enabled, the Ethernet Routing Switch can determine if an incoming unicast frame is a directed broadcast for one of its interfaces. If so, the switch forwards the datagram onto the appropriate network using a link-layer broadcast.

With IP directed broadcasting enabled on a VLAN, the Ethernet Routing Switch forwards direct broadcast packets in the following two ways:

- through a connected VLAN subnet to another connected VLAN subnet
- through a remote VLAN subnet to the connected VLAN subnet By default, this feature is disabled.

ARP

The Address Resolution Protocol (ARP) allows the Ethernet Routing Switch to dynamically learn Layer 2 Media Access Control (MAC) addresses, and to build a table with corresponding Layer 3 IP addresses.

Network stations using the IP protocol need both a physical (MAC) address and an IP address to transmit a packet. If a network station knows only the IP address of a network host, ARP enables the network station to determine the physical address of the network host and bind the 32-bit IP address to a 48-bit MAC address. A network station can use ARP across a single network only, and the network hardware must support physical broadcasts.

If a network station wants to send a packet to a host but knows only the host IP address, the network station uses ARP to determine the physical address of the host as follows:

1. The network station broadcasts a special packet, called an ARP request, that asks the host at the specified IP address to respond with its physical address.
2. All network hosts receive the broadcast message.
3. Only the specified host responds with its hardware address.
4. The network station then maps the host IP address to its physical address and saves the results in an address resolution table for future use.
5. The network station ARP table displays the association of the known MAC addresses to IP addresses.

The lifetime for the learned MAC addresses is a configurable parameter. The switch executes ARP lookups after this timer expires.

The default timeout value for ARP entries is 6 hours.

Static ARP

In addition to the dynamic ARP mechanism, the Ethernet Routing Switch supports a static mechanism that allows for static ARP entries to be added. With Static ARP, you can manually associate a device MAC address to an IP address. You can add and delete individual static ARP entries on the switch.

Proxy ARP

Proxy ARP allows the Ethernet Routing Switch to respond to an ARP request from a locally attached host that is intended for a remote destination. It does so by sending an ARP response back to the local host with the MAC address of the switch interface that is connected to the host subnet. The reply is generated only if the switch has an active route to the destination network.

With Proxy ARP enabled, the connected host can reach remote subnets without the need to configure default gateways.

The following figure is an example of proxy ARP operation. In this example, host B wants to send traffic to host C, so host B sends an ARP request for host C. However, the Avaya Ethernet Routing Switch 3500 Series, is between the two hosts, so the ARP message does not reach host C. To enable communication between the two hosts, the Avaya Ethernet Routing Switch 3500 Series, intercepts the message and responds to the ARP request with the IP address of host C but with the MAC address of the switch itself. Host B then updates its ARP table with the received information.

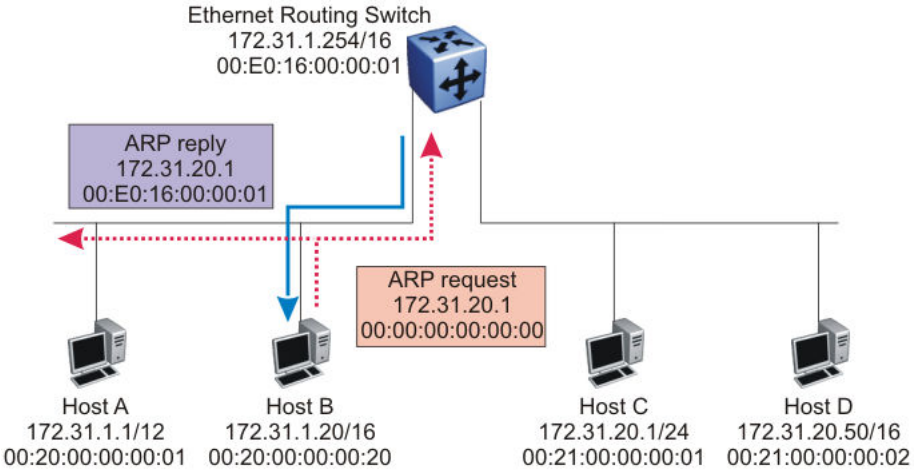

Figure 8: Proxy ARP Operation

Avaya recommends Proxy ARP as a temporary fix only, for example, if you are gradually moving hosts from one addressing scheme to another and you still want to maintain connectivity between the disparately-addressed devices. You do not want Proxy ARP running as a general rule because it causes hosts to generate ARP messages for every address that they want to reach on the Internet.

Routing feature capabilities and limitations

The following list describes the routing feature capabilities and limitations on the Ethernet Routing Switch:

- A maximum of 32local routes, and up to 32 static routes including the default route (Destination = 0.0.0.0 Mask = 0.0.0.0) are supported.
- The maximum number of management routes is 4.
- The maximum number of dynamic ARP entries is 480.
- The maximum number of static ARP entries is 256.

- When adding a static ARP entry for a VLAN subnet, the IP address associated with the MAC address must be in the subnet for the VLAN. Otherwise the following error message is returned:

  ```
  % Cannot modify settings

  IP address does not match with VLAN subnet
  ```

- UDP broadcast forwarding supports the following capabilities:
  - You can configure a maximum of 128 UDP port/protocol entries.
  - You can configure a maximum of 128 UDP forwarding lists.
  - You can configure a maximum of 16 ports (with their IP addresses) in one forwarding list.
  - You can bind a maximum of 16 VLANs to the same UDP forwarding list.

Chapter 5: IGMP fundamentals

This chapter provides an overview of IP multicast and Internet Group Management Protocol (IGMP). To support multicast traffic, the Avaya Ethernet Routing Switch 3500 Series, provides support for IGMP snooping.

Overview of IP multicast

Most traditional network applications such as Web browsers and e-mail employ unicast connections in which each client sets up a separate connection to a server to access specific data. However, with certain applications such as audio and video streaming, more than one client accesses the same data at the same time. With these applications, if the server sends the same data to each individual client using unicast connections, the multiple connections waste both server and network capacity. For example, if a server offers a 1 Mbit/sec live video stream for each client, a 100 Mbit/sec network interface card (NIC) on the server could be completely saturated after 90 client connections. The following figure shows an example of this waste of resources.

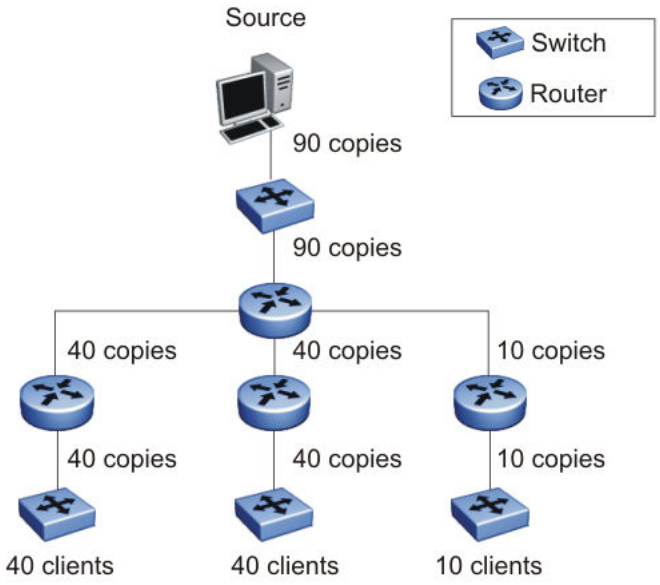

Figure 9: Wasteful propagation of multiple copies of the same unicast stream

Multicasting provides the ability to transmit only one stream of data to all the interested clients at the same time. The following figure shows a simple example of how multicasting works. The source of the multicast data forwards only one stream to the nearest downstream router, and

each subsequent downstream router forwards a copy of the same data stream to the recipients who are registered to receive it.

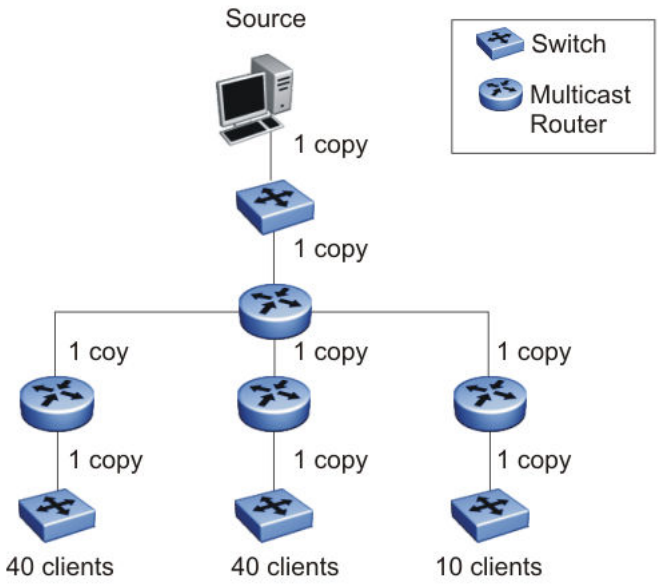

Figure 10: One stream replicated using multicasting

This one-to-many delivery mechanism is similar to broadcasting except that, while broadcasting transmits to all hosts in a network, multicasting transmits only to registered host groups. Because multicast applications transmit only one stream of data, which is then replicated to many receivers, multicasting saves a considerable amount of bandwidth.

Clients that want to receive the stream must register with the nearest multicast router to become a part of the receiving multicast group.

One downside to multicasting is that the multicast streams transmit data using User Datagram Protocol (UDP) packets, which are not as reliable as Transmission Control Protocol (TCP) packets.

Applications that use multicasting to transmit data include the following:

- multimedia conferencing
- real-time data multicasts (such as stock tickers)
- gaming and simulations

Multicast groups

To receive a multicast stream from a particular source, hosts must register with the nearest multicast router. The router adds all interested hosts to a multicast group, which is identified by a multicast IP address.

Multicast routers use Internet Group Membership Protocol (IGMP) to learn the existence of host group members on their directly attached subnets. To identify the hosts that want to be

IGMP fundamentals added to a group, a querier router sends out IGMP queries to each local network. A host that wants to belong to the group sends a response in the form of an IGMP membership report.

Each multicast router maintains a multicast routing table that lists each source, group (S,G) pair, which identifies the IP address of the source and the multicast address of the receiving group. For each (S,G) pair, the router maintains a list of downstream forwarding ports to which the multicast traffic is forwarded, and the upstream port where the multicast traffic is received.

Multicast addresses

Each multicast host group is assigned a unique multicast address. To reach all members of the group, a sender uses the multicast address as the destination address of the datagram.

An IP version 4 multicast address is a Class D address (the high-order bits are set to 1110) from 224.0.1.0 to 239.255.255.255. These addresses are assigned statically for use by permanent groups and dynamically for use by transient groups.

On the Ethernet Routing Switch 3500 Series, you cannot use 24-bit subnets like 224.0.0.0/24 and 224.128.0.0/24 for multicast data traffic. This restriction applies to the entire multicast address range from 224.0.0.0/8 to 239.128.0.0/8.

IGMP overview

IGMP is the Layer 3 protocol used by IP multicast routers to learn the existence of multicast group members on their directly attached subnets (see RFC 2236). With IGMP, hosts can register their desired group memberships to their local querier router. A multicast querier router communicates with hosts on a local network by sending IGMP queries. The router periodically sends a general query message to each local network of the router.

A host that wants to join a multicast group sends a response in the form of a membership report requesting registration with a group. After the querier router registers hosts to a group, it forwards all incoming multicast group packets to the registered host networks. As long as any host on a subnet continues to participate in the group, all hosts, including nonparticipating end stations on that subnet, receive the IP Multicast stream.

IGMP versions are backward compatible and can all exist together on a multicast network.

The following sections provide more details about the differences between the different IGMP versions.

IGMPv1 operation

IGMP version 1 is the simplest of the IGMP versions and is widely deployed.

IGMPv1 supports the following two message types:

- 0x11 – Membership Query message. Packets are sent to the all-systems multicast group (224.0.0.1).
- 0x12 – Membership Report message. Packets are sent to the group that the host intends to join.

The IGMPv1 router periodically sends host membership queries (also known as general queries) to its attached local subnets to inquire if any hosts are interested in joining any multicast groups. The interval between queries is a configurable value on the router. A host that wants to join a multicast group sends a membership report message to the nearest router, one report for each joined multicast group. After receiving the report, the router adds the Multicast IP address and the host port to its forwarding table. The router then forwards any multicast traffic for that multicast IP address to all member ports.

The router keeps a list of multicast group memberships for each attached network, and a Group Membership Interval timer for each membership. Repeated IGMP membership reports refresh the timer. If no reports are received before the timer expires, the router sends a query message.

In some cases, the host does not wait for a query before it sends report messages to the router. Upon initialization, the host can immediately issue a report for each of the multicast groups that it supports. The router accepts and processes these asynchronous reports the same way it accepts requested reports.

IGMPv1 leave process

After hosts and routers are in a steady state, they communicate in a way that minimizes the exchange of queries and reports. The designated routers set up a path between the IP Multicast stream source and the end stations, and periodically query the end stations to determine whether they want to continue to participate. As long as any host on the subnet continues to participate, all hosts, including nonparticipating end stations on the subnet, receive the IP Multicast stream.

If all hosts on the subnet leave the group, the router continues to send general queries to the subnet. If no hosts send reports after three consecutive queries, the router determines that no group members are present on the subnet.

IGMPv2 operation

IGMPv2 extends the IGMPv1 features by implementing a host leave message to quickly report group membership termination to the routing protocol. Instead of routers sending multiple

queries before determining that hosts have left a group, the hosts can send a leave message. This feature is important for multicast groups with highly volatile group membership.

The IGMPv2 join process is similar to the IGMPv1 join process.

IGMPv2 also implements a querier election process.

IGMPv2 adds support for the following three new message types:

- 0x11 – General Query and Group Specific Query message.
- 0x16 – Version 2 Membership Report (sent to the destination IP address of the group being reported)
- 0x17 – Version 2 Membership Leave message (sent to all-router [224.0.0.2] multicast address)

IGMPv2 also supports IGMPv1 messages.

Host leave process

With IGMPv2, if the host that issued the most recent report leaves a group, the host issues a leave message. The multicast router on the network then issues a group-specific query to determine whether other group members are present on the network. In the group-specific query message, the Group Address field is the group being queried (the Group Address field is 0 for the General Query message). If no host responds to the query, the router determines that no members belonging to that group exist on that interface.

The following figure shows an example of how IGMPv2 works.

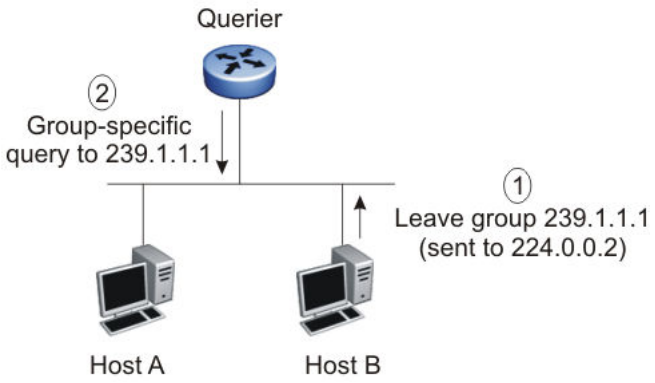

Figure 11: IGMPv2

In this example, the following occurs:

- The host sends a leave message (to 224.0.0.2).
- The router sends a group-specific query to group 239.1.1.1.

- No IGMP report is received.
- Group 239.1.1.1 times out.

Querier election process

Normally only one querier exists for each subnet. When multiple IGMPv2 routers are present on a network, the router with the lowest IP address is elected to send queries. All multicast routers start up as a querier on each attached network. If a multicast router receives a query message from a router with a lower IP address, the router with the higher IP address becomes a nonquerier on that network.

IGMPv3 operation

IGMPv3 adds support for source filtering. The IGMPv3 host can report its interest in receiving multicast packets from only specific source addresses, or the host can report its interest in receiving multicast packets from all but specific source addresses.

IGMPv3 is mostly used in voice and video conferences where multiple people can be part of the same conference. The IGMPv3 packet format adds a v3 Report message type (0x22) and includes Source-and-Group-specific Query messages.

The message type for Source-and-Group-specific Query message is 0x11, the same as IGMPv1 and IGMPv2. The different Query message versions are identified as follows:

- If the size of the IGMP message type is 8, then it is a v1 or v2 Query message.
- If the Group Address field is 0, then it is a General Query.
- If the Group Address field is a valid multicast IP address, then it is a Group-specific Query.
- If the Group Address field is a valid address and the Number of Sources field is nonzero, then it is a Group-and-Source specific Query message.

Each IGMPv3 Report contains a list of group records. The Group Record contains the multicast group address and the list of source addresses. The record type field specifies whether to INCLUDE or EXCLUDE the list of source addresses that are provided in the Source Address field. For example, to include packets from source 10.10.10.1, the report contains an INCLUDE(10.10.10.1) record.

The list of source addresses can be empty, which is represented by braces ({}), which means either to INCLUDE or EXCLUDE none. For example, the host that wants to receive packets from all group members can send a report with an EXCLUDE({}) record and a host that wants to leave a group can send a report with an INCLUDE({}) record, which is similar to a leave message.

In the following figure, hosts A, B, C, D, E, and F are part of a conference group G1. All hosts except F send a report for group G1 with the mode as INCLUDE(A, B, C, D, E, F) containing

all the source addresses. Host F, which is not interested in listening to C and D, sends a report to group G1 with the mode as EXCLUDE(C, D).

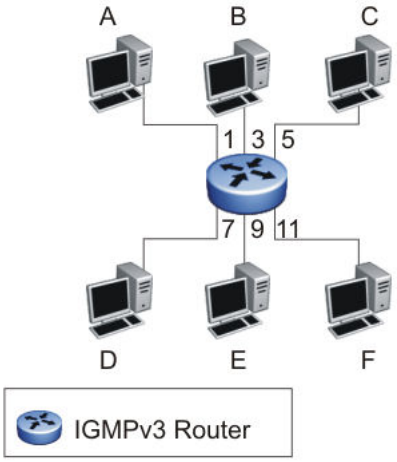

Figure 12: IGMPv3

The router adds the multicast IP address and the list of sources in the forwarding table. The router forwards the packets from A, B, E, and F to all ports. If the packets are received from C and D, it is forwarded to all ports except port 11.

IGMP requests for comment

For additional information about IGMP, see the following requests for comment (RFC):

- For IGMPv1, see RFC 1112.
- For IGMPv2, see RFC 2236.
- For IGMPv3, see RFC 3376
- For IGMP snooping, see RFC 4541.
- For IGMP management information bases (MIB), see RFC 2933.

IGMP snooping

If at least one host on a VLAN specifies that it is a member of a group, by default, the Avaya Ethernet Routing Switch 3500 Series, forwards to that VLAN all datagrams bearing the multicast address of that group. All ports on the VLAN receive the traffic for that group.

The following figure shows an example of this scenario. Here, the IGMP source provides an IP Multicast stream to a designated router. Because the local network contains receivers, the designated router forwards the IP Multicast stream to the network. Switches without IGMP snoop enabled flood the IP Multicast traffic to all segments on the local subnet. The receivers

requesting the traffic receive the desired stream, but so do all other hosts on the network. Although the nonparticipating end stations can filter the IP Multicast traffic, the IP Multicast traffic still exists on the subnet and consumes bandwidth.

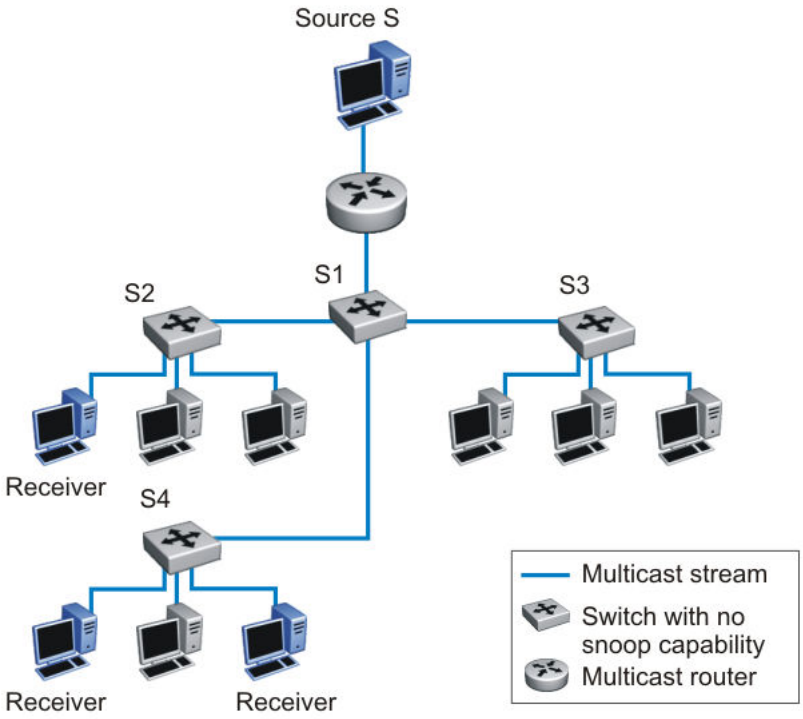

Figure 13: IP multicast propagation on a LAN without IGMP snooping

To prune ports that are not group members from receiving the group data, the Avaya Ethernet Routing Switch 3500 Series supports IGMP snoop for IGMPv1, IGMPv2, and IGMPv3. With IGMP snoop enabled on a VLAN, the switch forwards the multicast group data to only those ports that are members of the group. When using IGMP snoop, VLANs can provide the same benefit as IP Multicast routers, but in the local area.

The Avaya Ethernet Routing Switch 3500 Series, identifies multicast group members by listening to IGMP packets (IGMP reports, leaves, and queries) from each port. The switch suppresses the reports by not forwarding them out to other VLAN ports, forcing the members to continuously send their own reports. The switch uses the information gathered from the reports to build a list of group members. After the group members are identified, the switch blocks the IP Multicast stream from exiting any port that does not connect to a group member, thus conserving bandwidth.

As shown in the following figure, after the switches learn which ports are requesting access to the IP Multicast stream, all other ports not responding to the queries are blocked from receiving the IP Multicast data.

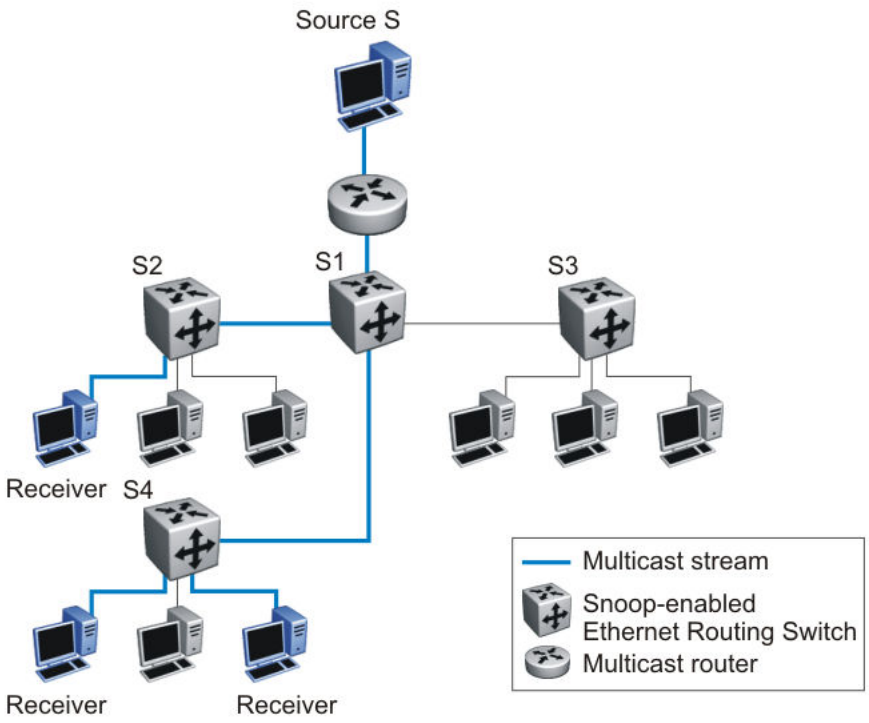

Figure 14: Ethernet Routing Switch running IGMP snooping

The switch continues to forward the IGMP membership reports from the hosts to the multicast routers, and forwards queries from multicast routers to all port members of the VLAN.

IGMPv3 snooping

In IGMPv3 snooping mode, the switch recognizes IGMPv3 reports and queries and can:

- recognize whether a source list is populated or blank
- identify the specific sources to filter
- understand and process all IGMPv3 record type

The following are supported:

- source filtering (INCLUDE, EXCLUDE, ALLOW, BLOCK of multicast sources)
- SSM (Source Specific Multicast)

The following table shows how IGMPv3 snooping handles different record types.

IGMP v3 record type	Without multicast source ({ })	Action	With multicast source(s) ({S1, S2...})	Action
MODE_IS_INCLUDE (1)	This is INCLUDE NONE.	LEAVE the group.	This is INCLUDE multicast sources.	JOIN the group. Discard multicast source information.
MODE_IS_EXCLUDE (2) .	This is EXCLUDE NONE.	JOIN the group.	This is EXCLUDE sources.	JOIN the group Discard multicast source information.
CHANGE_TO_INCLUDE_MODE (3)	This is include filter mode for multicast group.	LEAVE the group	This is include filter mode for multicast group.	JOIN the group. Discard multicast source information.
CHANGE_TO_EXCLUDE_MODE (4)	This is exclude filter mode for multicast group.	JOIN the group.	This is exclude filter mode for multicast group.	JOIN the group. Discard multicast source information.
ALLOW_NEW_SOURCES (5)	This type is for allowing new sources. This record type comes with sources. (This case may not happen.)	JOIN the group.	This type is for allowing new sources.	JOIN the group. Discard multicast source information.
BLOCK_OLD_SOURCES (6)	This type is for blocking existing sources.	JOIN the group.	This type is for blocking existing sources.	LEAVE the group. Discard multicast source information.

IGMP proxy

With IGMP snoop enabled, the switch can receive multiple reports for the same multicast group. Rather than forward each report upstream, the Ethernet Routing Switch 3500 Series can consolidate these multiple reports by using the IGMP proxy feature. With IGMP proxy enabled, if the switch receives multiple reports for the same multicast group, it does not transmit each report to the upstream multicast router. Instead, the switch forwards the first report to the querier and suppresses the rest. If new information emerges that another multicast group is

added or that a query is received because the last report is transmitted upstream, the report is then forwarded to the multicast router ports.

To enable IGMP Proxy, you must first activate IGMP snooping.

In the figure that follows, switches S1 to S4 represent a local area network (LAN) connected to an IP Multicast router. The router periodically sends Host Membership Queries to the LAN and listens for a response from end stations. All of the clients connected to switches S1 to S4 are aware of the queries from the router.

One client, connected to S2, responds with a host membership report. Switch S2 intercepts the report from that port, and generates a proxy report to its upstream neighbor, S1. Also, two clients connected to S4 respond with host membership reports, causing S4 to intercept the reports and to generate a consolidated proxy report to its upstream neighbor, S1.

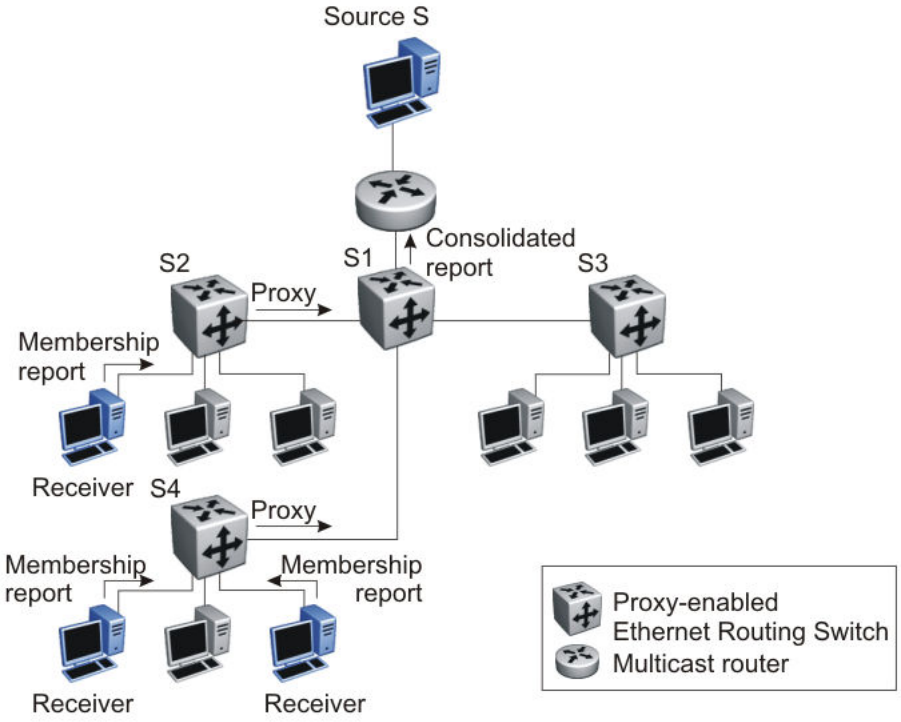

Figure 15: Ethernet Routing Switch running IGMP proxy

Switch S1 treats the consolidated proxy reports from S2 and S4 as if they were reports from any client connected to its ports, and generates a consolidated proxy report to the designated router. In this scenario, the router receives a single consolidated report from that entire subnet.

The consolidated proxy report generated by the switch remains transparent to Layer 3 of the International Standardization Organization, Open Systems Interconnection (ISO/OSI) model. (The switch IP address and Media Access Control [MAC] address are not part of proxy report generation.) The last reporting IGMP group member in each VLAN represents all of the hosts in that VLAN and IGMP group.

IGMPv3 proxy

With IGMPv3 proxy enabled, if the switch receives multiple reports for the same multicast group, it does not transmit each report to the upstream multicast router. Instead, the switch forwards the first report to the querier and suppresses the rest.

If new information emerges, for example if the switch adds another multicast group or receives a query since the last report was transmitted upstream, then the switch forwards a new report to the multicast router ports.

Forwarding of reports

When forwarding IGMP membership reports from group members, the Ethernet Routing Switch 3500 Series forwards the reports only to those ports where multicast routers are attached. To do this, the switch maintains a list of multicast querier routers and the multicast router (mrouter) ports on which they are attached. The switch learns of the multicast querier routers by listening to the queries sent by the routers where source address is not 0.0.0.0.

Static mrouter port and nonquerier

If two IGMP routers are active on a VLAN, the router with the lower IP address is the querier, and the router with the higher IP address operates as a nonquerier. Only querier routers forward IGMP queries on the VLAN; nonqueriers do not forward IGMP queries. IGMP snoop considers the port on which the IGMP query is received as the active IGMP multicast router (mrouter) port. IGMP snoop is not aware of nonquerier IGMP routers.

By default, IGMP snoop forwards reports to the IGMP querier router only. To allow the switch to forward reports to the nonquerier router as well, you can configure the port connected to the nonquerier as a static mrouter port.

The following figure shows how static mrouter ports operate. In this case, the Ethernet Routing Switch 3500 Series has port members 5/1 and 6/1 connected to IGMP routers in VLAN 10. Router 1 is the IGMP querier because it has a lower IP address than router 2. Router 2 is then considered the nonquerier.

By default, the switch learns of the multicast querier routers by listening to the IGMP queries. In this case, port 6/1 connected to querier router 1 is identified as an mrouter port.

To forward reports to IGMP router 2 as well, you can configure port 5/1 on the switch as a static mrouter port. In this case, the IGMP reports are forwarded to both routers.

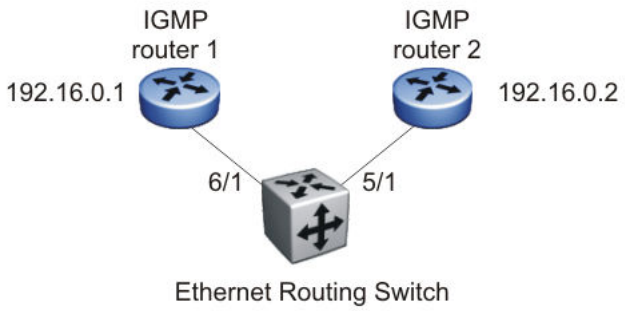

Figure 16: Static mrouter port and nonquerier

Unknown multicast packet filtering

With IGMP snoop enabled, if the switch receives multicast packets with destination addresses that it has not already registered using IGMP reports, the switch floods all such packets to all ports on the VLAN. All unknown multicast streams of a group are flooded on the VLAN until at least one port in the VLAN becomes a member of that group.

On the switch, you can enable the unknown multicast filtering feature so that the unknown multicast packets are not flooded on the VLAN. To enable unknown multicast filtering, you can use the `vlan igmp unknown-mcast-no-flood ACLI` command.

With this feature enabled, the switch forwards all unknown multicast traffic to IGMP static mrouter ports only. The traffic is not forwarded to dynamically discovered mrouter ports.

If you require unknown multicast traffic to be forwarded to certain ports (for example, to forward Layer 3 multicast routing traffic), set the ports as static mrouter ports. Avaya recommends that you enable this feature after IGMP snooping is enabled. User settings for the unknown multicast filtering feature are stored in NVRAM.

Allowing a multicast MAC address to flood all VLANs The unknown multicast filtering feature introduces a potential problem after a Layer 2 VLAN is placed between two Layer 3 switches that are exchanging protocol packets such as OSPF. Since the protocols do not join a multicast group, the associated MAC addresses cannot be identified by the IGMP snooping process. These packets are dropped by the Layer 2 switch because the unknown multicast filtering feature is enabled. The two Layer 3 switches can never establish adjacencies and the OSPF protocol fails.

Using the `vlan igmp unknown-mcast-allow-flood ACLI` command, you can specify MAC addresses or multicast IP addresses that need to be flooded on the switch even when the unknown multicast filtering feature is enabled. The specified MAC or IP addresses are added to the allow-flood table for all VLANs. Any matching packets are flooded on all ports of a VLAN.

Robustness value

As part of the IGMP snooping configuration, use the robustness value to configure the switch to offset expected packet loss on a subnet. If you expect a network to lose query packets, increase the robustness value.

This value is equal to the number of expected query packet losses for each query interval, plus 1. The range is from 2 to 255, and the default is 2. The default value of 2 means that one query for each query interval can be dropped without the querier aging out.

IGMP snooping configuration rules

The IGMP snooping feature operates according to specific configuration rules. When configuring your switch for IGMP snooping, consider the following rules that determine how the configuration reacts in any network topology:

- The switch supports up to 59 multicast groups.

  If the multicast group table reaches its limit, a new entry cannot be added with a JOIN message or a new sender identifying a new group. The multicast stream from the new sender is discarded by the hardware. New entries can be added again when the table is not full.

- You cannot configure port mirroring on a static mrouter port.
- If you configure a Multi-Link Trunk member as a static mrouter port, all the Multi-Link Trunk members become static mrouter ports. Also, if you remove a static mrouter port that is a Multi-Link Trunk member, all Multi-Link Trunk members are automatically removed as static mrouter port members.
- When you specify MAC or IP addresses to be flooded on the switch, the specified addresses are flooded only on the VLAN specified within the ACLI command. This way, you can flood MAC or IP addresses for specific VLANs only.
- Static mrouter ports must be port members of at least one VLAN.
- If you configure a port as a static mrouter port, it is configured as a static mrouter port for all VLANs on that port. The IGMP configuration is propagated through all VLANs of that port.
- If you remove a static mrouter port, the membership for that port is removed from all VLANs of that port.
- When Spanning Tree is enabled, the switch learns IGMP groups only on ports that are not in Listening or Blocking Spanning Tree states (or, when in RSTP/MSTP mode, only on ports that are in the Designated state). The switch also learns the groups if STP is disabled on a port.
- The IGMP snooping feature is not Rate Limiting-dependent.

- You must enable the IGMP snooping feature before you can enable the IGMP proxy feature.
- You can specify static mrouter ports per VLAN.

Important:

Because IGMP snooping is set up per VLAN, all IGMP changes are implemented according to the VLAN configuration for the specified ports.

Default IGMP values

The following table lists the default IGMP values on the Ethernet Routing Switch.

Parameters	Range	Default Value
Snooping	Enable/Disable	Disable
Version	1-3	2
Proxy	Enable/Disable	Disable
Query Interval	0-65535	125
Robustness Value	2-255	2

IGMP snooping interworking with Windows clients

This section describes an interworking issue between Windows clients and the Ethernet Routing Switches when IGMP snoop is enabled for multicast traffic.

Under normal IGMP snoop operation, as soon as a client joins a specific multicast group, the group is no longer unknown to the switch, and the switch sends the multicast stream only to the ports which request it.

To force a Windows client to only use IGMPv1 or IGMPv2 reports, change the TCP/IP settings in the Windows Registry located under the following registry key:

Note:

ERS3500 Release 5.0 now supports IGMPv3, and therefore, these settings are only required if you are using IGMPv1, or IGMPv2.

```
HKEY_LOCAL_MACHINE
\SYSTEM
\CurrentControlSet
\Services
\Tcpip
\Parameters
```

The specific parameter which controls the IGMP Version is:

```
IGMPVersion
Key: Tcpip\Parameters
Value Type: REG_DWORD—Number
Valid Range: 2, 3, 4
Default: 4
```

To set the Windows Client to only utilize IGMPv2, change the IGMPVersion parameter to 3 (2 specifies IGMPv1, 3 specifies IGMPv2, and 4 specifies IGMPv3).

The IGMPVersion parameter may not be present in the list of the TCP/IP parameters. By default, the system assumes the IGMPv3 value (4). To configure the system for IGMPv2, create the parameter as a DWORD key in the registry and specify Decimal 3.

Important:

If you edit the Windows registry incorrectly, you can severely damage your system. As a minimal safeguard, back up your system data before undertaking changes to the registry.

Chapter 6: IP routing configuration using ACLI

This chapter describes the procedures you can use to configure routable VLANs using the ACLI.

The Avaya Ethernet Routing Switch 3500 Series, are Layer 3 switches. This means that a regular Layer 2 VLAN becomes a routable Layer 3 VLAN if an IP address is attached to the VLAN. When routing is enabled in Layer 3 mode, every Layer 3 VLAN is capable of routing and carrying the management traffic. You can use any Layer 3 VLAN instead of the Management VLAN to manage the switch.

For more information about creating and configuring VLANs, see *Configuration—VLANs, Spanning Tree, and Link Aggregation (NN47215-501)*.

IP routing configuration procedures

To configure inter-VLAN routing on the switch, perform the following steps:

Procedure

1. Enable IP routing globally.
2. Assign IP addresses to multiple VLANs.

   Routing is automatically enabled on the VLAN after you assign an IP address to it.

   In the preceding procedure, you are not required to enable IP routing as the first step. You can configure all IP routing parameters on the Avaya Ethernet Routing Switch 3500 Series, before you enable routing on the switch.

Configuring global IP routing status using ACLI

Use this procedure to enable and disable global routing at the switch level. By default, routing is disabled.

Procedure

1. Log on to ACLI in Global Configuration command mode.

2. At the command prompt, enter the following command:
   `[no] ip routing`

Variable definitions

The following table describes the parameters for the **`ip routing`**command.

Variable	Value
no	Disables IP routing on the switch.

Displaying global IP routing status using ACLI

Use this procedure to display the status of IP routing on the switch.

Procedure

1. Log on to ACLI in User EXEC command mode.
2. At the command prompt, enter the following command:
   `show ip routing`

Configuring an IP address for a VLAN using ACLI

To enable routing on a VLAN, you must first configure an IP address on the VLAN.

Procedure

1. Log on to ACLI in VLAN Interface Configuration command mode.
2. At the command prompt, enter the following command:
   `[no] ip address <ipaddr> <mask> [<MAC-offset>]`

Variable definitions

The following table describes the parameters for the **ip address** command.

Variable	Value
[no]	Removes the configured IP address and disables routing on the VLAN.
<ipaddr>	Specifies the IP address to attach to the VLAN.
<mask>	Specifies the subnet mask to attach to the VLAN.
[<MAC-offset>]	Specifies the value used to calculate the VLAN MAC address, which is offset from the switch MAC address. RANGE: The valid range is 1-256. Specify the value 1 for the Management VLAN only. If no MAC offset is specified, the switch applies one automatically.

Configuring IP routing status on a VLAN using ACLI

Use this procedure to enable and disable routing for a particular VLAN.

Procedure

1. Log on to ACLI in VLAN Interface Configuration command mode.
2. At the command prompt, enter the following command:
   ```
   [default] [no] ip routing
   ```

Variable definitions

The following table describes the parameters for the **ip routing** command.

Variable	Value
default	Disables IP routing on the VLAN.

Variable	Value
no	Disables IP routing on the VLAN.

Displaying the IP address configuration and routing status for a VLAN using ACLI

Use this procedure to display the IP address configuration and the status of routing on a VLAN.

Procedure

1. Log on to ACLI in Privileged Exec command mode.
2. At the command prompt, enter the following command:
   `show vlan ip [vid <vid>]`

   The following information is displayed:
   - Vid — Specifies the VLAN ID
   - ifIndex — Specifies an index entry for the interface
   - Address — Specifies the IP address associated with the VLAN
   - Mask — Specifies the mask
   - MacAddress — Specifies the MAC address associated with the VLAN
   - Offset — Specifies the value used to calculate the VLAN MAC address, which is offset from the switch MAC address
   - Routing — Specifies the status of routing on the VLAN: enabled or disabled

Variable definitions

The following table describes the parameters for the **`show vlan ip`** command.

Variable	Value
[vid<vid>]	Specifies the VLAN ID of the VLAN to be displayed. RANGE: 1–4094.

Displaying IP routes using ACLI

Use this procedure to display all active routes on the switch.

Procedure

1. Log on to ACLI in User EXEC command mode.
2. At the command prompt, enter the following command:
   `show ip route [<dest-ip>] [-s <subnet> <mask>]`

   The following information is displayed:

   - DST — Identifies the route destination
   - MASK — Identifies the route mask
   - NEXT — Identifies the next hop in the route
   - COST — Identifies the route cost
   - VLAN — Identifies the VLAN ID on the route
   - PORT — Specifies the ports
   - PROT — Specifies the routing protocols. Options are LOC (local route) or STAT (static route)
   - TYPE — Indicates the type of route as described by the Type Legend
   - PRF — Specifies the route preference

Variable definitions

The following table describes the parameters for the **`show ip route`** command.

Variable	Value
<dest-ip>	Specifies the destination IP address of the routes to display.
[-s<subnet><mask>]	Specifies the destination subnet of the routes to display.

Chapter 7: Static route configuration using ACLI

This chapter describes the procedures you can use to configure static routes using the ACLI.

Configuring a static route using ACLI

Create static routes to manually configure a path to destination IP address prefixes.

Before you begin

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the VLANs to be routed.

Procedure

1. Log on to ACLI in Global Configuration command mode.
2. At the command prompt, enter the following command:

```
[no] ip route <dest-ip> <mask> <next-hop> [<cost>] [disable]
[enable] [weight <cost>]
```

Variable definitions

The following table describes the parameters for the **`ip route`** command.

Variable	Value
[no]	Removes the specified static route.
<dest-ip>	Specifies the destination IP address for the route being added. DEFAULT: 0.0.0.0 is considered the default route.
<mask>	Specifies the destination subnet mask for the route being added.

Variable	Value
<next-hop>	Specifies the next hop IP address for the route being added.
[<cost>]	Specifies the weight, or cost, of the route being added. RANGE: 1–65535
[enable]	Enables the specified static route.
[disable]	Disables the specified static route.
[weight<cost>]	Changes the weight, or cost, of an existing static route. RANGE: 1–65535

Displaying static routes using ACLI

Use this procedure to display all static routes, whether these routes are active or inactive.

Procedure

1. Log on to ACLI in User EXEC command mode.
2. At the command prompt, enter the following command:
   `show ip route static [<dest-ip>] [-s <subnet> <mask>]`

   The following information is displayed:

   - DST — Identifies the route destination
   - MASK — Identifies the route mask.
   - NEXT — Identifies the next hop in the route.
   - COST — Identifies the route cost.
   - PREF — Specifies the route preference.
   - LCLNHOP — Specifies the local next hop status.
   - STATUS — Specifies the static route status. Options are ACTIVE (in use and present in routing table) or INACTV (not in use and not present in routing table).
   - ENABLE — Specifies the administrative state of the static route. Options are TRUE (administratively enabled) or FALSE (administratively disabled).

Variable definitions

The following table describes the parameters for the **show ip route static** command.

Variable	Value
<dest-ip>	Specifies the destination IP address of the static routes to display.
[-s <subnet> <mask>]	Specifies the destination subnet of the routes to display.

Configuring a management route using ACLI

Use this procedure to create a management route to the far end network, with a next-hop IP address from the management VLAN's subnet. You can configure a maximum of four management routes on the switch.

Before you begin

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the management VLAN interface.

Procedure

1. Log on to ACLI in Global Configuration command mode.
2. At the command prompt, enter the following command:
   ```
   [no] ip mgmt route <dest-ip> <mask> <next-hop>
   ```

Variable definitions

The following table describes the parameters for the **ip mgmt route** command.

Variable	Value
[no]	Removes the specified management route.
<dest-ip>	Specifies the destination IP address for the route being added.
<mask>	Specifies the destination subnet mask for the route being added.

Variable	Value
<next-hop>	Specifies the next hop IP address for the route being added.

Displaying the management routes using ACLI

Use this procedure to display the static routes configured for the management VLAN.

Procedure

1. Log on to ACLI in User EXEC command mode.
2. At the command prompt, enter the following command:
   `show ip mgmt route`

   The following information is displayed:

   - Destination IP — Identifies the route destination.
   - Subnet Mask — Identifies the route mask.
   - Gateway IP — Identifies the next hop in the route.

Chapter 8: DHCP relay configuration using ACLI

This chapter describes the procedures you can use to configure Dynamic Host Configuration Protocol (DHCP) relay using the ACLI.

Important:

DHCP relay uses a hardware resource that is shared by switch Quality of Service applications. When DHCP relay is enabled globally, the Quality of Service filter manager will not be able to use precedence 3 for configurations. For the filter manager to be able to use this resource, DHCP relay must be disabled for the entire unit.

Prerequisites to DHCP relay configuration using ACLI

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the VLAN to be set as the DHCP relay agent.
- Ensure that a route (local or static) to the destination DHCP server is available on the switch.

DHCP relay configuration procedures using ACLI

Use the following procedure to configure DHCP relay.

Procedure

1. Ensure that DHCP relay is enabled globally. (DHCP relay is enabled by default).
2. Configure the DHCP relay forwarding path by specifying a local VLAN as the DHCP relay agent and the remote DHCP server as the destination.
3. Enable DHCP relay for the specific VLAN.

Enabling or disabling global DHCP relay using ACLI

Use the following procedure to enable or disable global DHCP relay. DHCP relay is enabled by default.

Procedure

1. Log on to ACLI in Global configuration command mode.
2. At the command prompt, enter the following command:
   `ip dhcp-relay` to enable

   OR

   `no ip dhcp-relay` to disable

Setting global DHCP relay to default using ACLI

Use the following procedure to set DHCP relay to default settings for the switch. DHCP relay is enabled by default.

Procedure

1. Log on to ACLI in Global configuration command mode.
2. At the command prompt, enter the following command:
   `default ip dhcp-relay`

Displaying the global DHCP relay status using ACLI

Use this procedure to display the current DHCP relay status for the switch.

Procedure

1. Log on to ACLI in User EXEC command mode.
2. At the command prompt, enter the following command:

```
show ip dhcp-relay
```

Variable definitions

The following table describes the parameters for the **`ip dhcp-relay`** command.

Variable	Value
default	Sets DHCP relay to default settings.
no	Disables DHCP relay.
show	Shows the status of the DHCP relay.

Displaying IP DHCP client parameters using ACLI

Use the following procedure to display IP DCHP client parameters for the switch.

Procedure

1. Log on to ACLI in Global configuration command mode.
2. At the command prompt, enter the following command:
   ```
   show ip dhcp client lease
   ```

Specifying a local DHCP relay agent and remote DHCP server using ACLI

Use this procedure to specify a local VLAN as a DHCP relay agent on the forwarding path to a remote DHCP server. The DHCP relay agent can forward DHCP client requests from the local network to the DHCP server in the remote network.

The DHCP relay feature is enabled by default, and the default mode is BootP-DHCP.

Procedure

1. Log on to ACLI in Global configuration command mode.
2. At the command prompt, enter the following command:

```
[no] ip dhcp-relay fwd-path <relay-agent-ip> <DHCP-server>
[enable] [disable] [mode {bootp | bootp-dhcp | dhcp}]
```

Variable definitions

The following table describes the parameters for the **`ip dhcp-relay fwd-path`** command.

Variable	Value
[no]	Removes the specified DHCP forwarding path.
<relay-agent-ip>	Specifies the IP address of the VLAN that serves as the local DHCP relay agent.
<DHCP-server>	Specifies the address of the remote DHCP server to which DHCP packets are to be relayed.
[enable]	Enables the specified DHCP relay forwarding path.
[disable]	Disables the specified DHCP relay forwarding path.
[mode {bootp | bootp-dhcp | dhcp}]	Specifies the DHCP relay mode: • BootP only • BootP and DHCP • DHCP only If you do not specify a mode, the default DHCP and BootP is used.

Displaying the DHCP relay configuration using ACLI

Use this procedure to display the current DHCP relay agent configuration.

Procedure

1. Log on to ACLI in User EXEC command mode.
2. At the command prompt, enter the following command:

```
show ip dhcp-relay fwd-path
```

Configuring DHCP relay on a VLAN using ACLI

Use this procedure to configure the DHCP relay parameters on a VLAN.

To enable DHCP relay on the VLAN, enter the command with no optional parameters.

Procedure

1. Log on to ACLI in VLAN Interface Configuration command mode.
2. At the command prompt, enter the following command:

```
[no] ip dhcp-relay [broadcast][clear counters][min-sec <min-sec>] [mode {bootp | dhcp | bootp_dhcp}][Option-82]
```

Variable definitions

The following table describes the parameters for the **`ip dhcp-relay`** command.

Variable	Value
[no]	Disables DHCP relay on the specified VLAN.
[broadcast]	Enables the broadcast of DHCP reply packets to the DHCP clients on this VLAN interface.
[Clear Counters]	Clear the existing number of counters and restart the counters.
min-sec<*min-sec*>	Indicates the min-sec value. The switch immediately forwards a BootP/DHCP packet if the secs field in the BootP/DHCP packet header is greater than the configured min-sec value; otherwise, the packet is dropped. RANGE: 0–65535 DEFAULT: The default is 0.

Variable	Value
mode {bootp | dhcp | bootp_dhcp}	Specifies the type of DHCP packets this VLAN supports: • bootp - Supports BootP only • dhcp - Supports DHCP only • bootp_dhcp - Supports both BootP and DHCP
[Option-82]	Specifies the DHCP Option 82 subscriber ID for the port.

Displaying the DHCP relay configuration for a VLAN using ACLI

Use this procedure to display the current DHCP relay parameters configured for a VLAN.

Procedure

1. Log on to ACLI in Privileged EXEC command mode.
2. At the command prompt, enter the following command:
   `show vlan dhcp-relay [<vid>]`

   The following information is displayed:

   - IfIndex — Indicates the VLAN interface index.
   - MIN_SEC — Indicates the min-sec value. The switch immediately forwards a bootP/DHCP packet if the secs field in the BootP/DHCP packet header is greater than the configured min-sec value; otherwise, the packet is dropped.
   - ENABLED — Indicates whether DHCP relay is enabled on the VLAN.
   - MODE — Indicates the type of DHCP packets this interface supports. Options include none, BootP, DHCP, and both.
   - ALWAYS_BROADCAST — Indicates whether DHCP reply packets are broadcast to the DHCP client on this VLAN interface.

Variable definitions

The following table describes the parameters for the **show vlan dhcp-relay** command.

Variable	Value
[<vid>]	Specifies the VLAN ID of the VLAN to be displayed. RANGE: 1–4094

Displaying DHCP relay counters using ACLI

Use this procedure to display the current DHCP relay counters. This includes the number of requests and the number of replies.

Procedure

1. Log on to ACLI in User EXEC command mode.
2. At the command prompt, enter the following command:
   `show ip dhcp-relay counters`

   The following information is displayed:

   - INTERFACE — Indicates the interface IP address of the DHCP relay agent.
   - REQUESTS — Indicates the number of DHCP requests.
   - REPLIES — Indicates the number of DHCP replies.

Clearing DHCP relay counters for a VLAN using ACLI

Use this procedure to clear the DHCP relay counters for a VLAN.

Procedure

1. Log on to ACLI in VLAN Interface Configuration command mode.
2. At the command prompt, enter the following command:

```
ip dhcp-relay clear-counters
```

Configuring DHCP Relay Option 82 globally using ACLI

To enable or disable the DHCP Relay Option 82 at the switch level, you can configure Option 82 for DHCP relay globally.

Procedure

1. Log on to ACLI in Global Configuration command mode.
2. At the command prompt, enter the following command:
   ```
   [no|default] ip dhcp-relay option82
   ```

Variable definitions

The following table describes the parameters for the **`ip dhcp-relay option82`** command.

Variable	Value
default	Resets DHCP Relay Option 82 to default values. DEFAULT: Default value is disabled.
no	Disables DHCP Relay Option 82 for the switch.

Configuring DHCP Relay with Option 82 for a VLAN using ACLI

Perform the following procedure to configure DHCP Relay with Option 82 for a VLAN.

Procedure

1. Log on to ACLI in Interface VLAN configuration command mode.
2. At the command prompt, enter the following command:

```
ip dhcp-relay option82
```

Configuring DHCP Forwarding Maximum Frame size using ACLI

You can specify the maximum frame size the DHCP relay agent can forward to the DHCP server. While the switch implementation permits configuration of the maximum DHCP packet size up to 1536 bytes, the default maximum size is 576 bytes.

Use the following procedure to configure DHCP Forwarding maximum frame size.

Procedure

1. Log on to ACLI in Global Configuration command mode.
2. At the command prompt, enter the following command:
   ```
   ip dhcp-relay max-frame <576–1536>
   ```

Assigning a DHCP Relay Option 82 subscriber ID to a port using ACLI

To associate an alphanumeric character string with the Option 82 function for a port, you can assign a DHCP Relay Option 82 subscriber ID to the port.

Procedure

1. Log on to ACLI in FastEthernet Interface command mode.
2. At the command prompt, enter the following command:
   ```
   [no|default] ip dhcp-relay option82-subscriber-id <Word 1–
   255>
   ```

Variable definitions

The following table describes the parameters for the **ip dhcp-relay option 82-subscriber-id** command.

Variable	Value
default	Resets DHCP Relay Option 82 subscriber ID to the default value. DEFAULT: The default is disabled.
no	Removes DHDP Relay Option 82 subscriber ID from a port.
Word	Specifies the DHCP Relay Option 82 subscriber ID for the port. The value is a character string between 1 and 255 characters.

Displaying DHCP Relay using ACLI

Use the following procedure to display the state of the DHCP Relay, DHCP Relay Option 82, and DHCP Relay maximum frame size.

Procedure

1. Log on to ACLI in Global Configuration command mode.
2. At the command prompt, enter the following command:
   `show ip dhcp-relay`

Example

```
3524GT-PWR+>enable
3524GT-PWR+#configure
Configuring from terminal or network [terminal]? terminal
Enter configuration commands, one per line.  End with CNTL/Z.
3524GT-PWR+(config)#show ip dhcp-relay
==============================
       DHCP Relay Global
==============================
DHCP relay is enabled
DHCP relay option82 is disabled
DHCP relay max-frame is 576
3524GT-PWR+(config)#
```

Chapter 9: UDP broadcast forwarding configuration using ACLI

This chapter describes the procedures you can use to configure UDP broadcast forwarding using ACLI. UDP broadcast forwarding is a general mechanism for selectively forwarding limited UDP broadcasts received on an IP interface to a configured IP address.

You cannot enable or disable the UDP broadcast forwarding feature on a global level. When you attach the first UDP forwarding list to a VLAN interface, the feature is enabled. When you remove the last UDP forwarding list from a VLAN, the feature is disabled.

Prerequisites to UDP broadcast forwarding using ACLI

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the VLAN to be configured as a UDP forwarding interface.
- Ensure that a route (local or static) to the destination address is available on the switch.

Important:

If you configure EAPOL on the switch, enable EAPOL prior to enabling UDP Forwarding, otherwise the UDP broadcast traffic matching UDP forward lists is forwarded regardless of the EAPOL port state (authorized, force unauthorized, or auto).

UDP broadcast forwarding configuration procedures

To configure UDP broadcast forwarding, perform the following steps:

1. Create UDP protocol entries that specify the protocol associated with each UDP port that you want to forward.
2. Create a UDP forwarding list that specifies the destination IP addresses for each forwarding UDP port. (You can create up to 128 UDP forwarding lists.)
3. Apply UDP forwarding lists to local VLAN interfaces.

Configuring UDP protocol table entries using ACLI

Use the following procedure to create UDP protocol table entries that identify the protocols associated with specific UDP ports to forward.

Procedure

1. Log on to ACLI in Global Configuration command mode.
2. At the command prompt, enter the following command:
   ```
   ip forward-protocol udp [<forwarding_port> <protocol_name>]
   ```

Variable definitions

The following table describes the parameters for the **`ip forward-protocol udp`** command.

Variable	Value
<forwarding_port>	Specifies the UDP port number. RANGE: 1–65535
<protocol_name>	Specifies the UDP protocol name.

Displaying the UDP protocol table using ACLI

Use the following procedure to display the configured UDP protocol table entries.

Procedure

1. Log on to ACLI in User Exec command mode.
2. At the command prompt, enter the following command:
   ```
   show ip forward-protocol udp
   ```
   The following information is displayed:
   - UDP_Port — Indicates the UDP ports.

• PROTOCOL_NAME — Indicates the name of the associated protocol.

Configuring a UDP forwarding list using ACLI

Use the following procedure to configure a UDP forwarding list, which associates UDP forwarding ports with destination IP addresses. Each forwarding list can contain multiple port/destination entries.

A maximum of 16 port/destination entries per forwarding list and up to 128 forwarding lists can be configured.

Procedure

1. Log on to ACLI in Global Configuration command mode.
2. At the command prompt, enter the following command:
   ```
   ip forward-protocol udp portfwdlist <forward_list> <udp_port>
   <dest_ip> [name <list_name>]
   ```

Variable definitions

The following table describes the parameters for the **`ip forward-protocol udp portfwdlist`** command.

Variable	Value
<forward_list>	Specifies the ID of the UDP forwarding list. RANGE: 1–128
<udp_port>	Specifies the port on which the UDP forwarding originates.
<dest_ip>	Specifies the destination IP address for the UDP port.
<list_name>	Specifies the name of the UDP forwarding list being created (maximum 15 characters).

Applying a UDP forwarding list to a VLAN using ACLI

Use the following procedure to associate a UDP forwarding list with a VLAN interface. One list can be associated at a time.

The same UDP forwarding list can be associated to a maximum of 16 different VLANs.

Note:

Due to hardware limitations, a forwarding list cannot be applied unless a QoS filter is free. To obtain a free QoS filter, you can disable DHCP Relay (if not used) or use the following ACLI commands:

```
3510GT-PWR+(config)#qos if-group name <name of the interface group> class
unrestricted
3510GT-PWR+(config)#qos if-assign port all name <name of the interface group>
```

Procedure

1. Log on to ACLI in VLAN Interface Configuration command mode.
2. At the command prompt, enter the following command:
   ```
   ip forward-protocol udp [vlan <vid>] [portfwdlist
   <forward_list>] [broadcastmask <bcast_mask>] [maxttl
   <max_ttl>]
   ```

Variable definitions

The following table describes the parameters for the **`ip forward-protocol udp`** command.

Variable	Value
<vid>	Specifies the VLAN ID on which to attach the UDP forwarding list. This parameter is optional, and if not specified, the UDP forwarding list is applied to the interface specified in the **`interface vlan`** command.
<forward_list>	Specifies the ID of the UDP forwarding list to attach to the selected VLAN interface.
<bcast_mask>	Specifies the 32-bit mask used by the selected VLAN interface to make forwarding

Variable	Value
	decisions based on the destination IP address of the incoming UDP broadcast traffic. If you do not specify a broadcast mask value, the switch uses the mask of the interface to which the forwarding list is attached
<max_ttl>	Specifies the time-to-live (TTL) value inserted in the IP headers of the forwarded UDP packets coming out of the selected VLAN interface. DEFAULT: 4

Note:

If you specify maxttl and/or broadcastmask values with no portfwdlist specified, the switch saves the settings for this interface. If you subsequently attach portfwdlist to this interface without defining the maxttl and/or broadcastmask values, the saved parameters are automatically attached to the list. But, if when specifying the portfwdlist, you also specify the maxttl and/or broadcastmask, your specified properties are used, regardless of any previous configurations.

Displaying the UDP broadcast forwarding configuration using ACLI

Use the following procedure to display the UDP broadcast forwarding configuration.

Procedure

1. Log on to ACLI in User Exec command mode.
2. At the command prompt, enter the following command:
   ```
   show ip forward-protocol udp [interface [vlan <1–4094>]]
   [portfwdlist [<portlist>]
   ```
   The following information is displayed:
   - UDP_PORT — Indicates the UDP ports.
   - PROTOCOL_NAME — Indicates the name of the protocol.

   The following information is displayed for the UDP interfaces command:
   - INTF_ADDR — Indicates the IP address of the interface.
   - FWD LISTID — Identifies the UDP forwarding policy.

- MAXTTL — Indicates the maximum TTL.
- RXPKTS — Indicates the number of received packets.
- FWDOKTS — Indicates the number of forwarded packets.
- DRPDEST UNREACH — Indicates the number of dropped packets that cannot reach the destination.
- DRP_UNKNOWN PROTOCOL — Indicates the number of packets dropped with an unknown protocol.
- BDCASTMASK — Indicates the value of the broadcast mask.

The following information is displayed for the UDP portfwdlist command:

- LIST_ID — Specifies the UDP forwarding policy number.
- NAME — Specifies the name of the UDP forwarding policy.

Variable definitions

The following table describes the parameters for the **`show ip forward-protocol udp`** command.

Variable	Value
[interface [vlan <1–4094>]]	Displays the configuration and statistics for a VLAN interface. If no VLAN is specified, the configuration for all UDP forwardingenabled VLANs is displayed.
[portfwdlist [<forward_list>]]	Displays the specified UDP forwarding list. If no list is specified, a summary of all forwarding lists is displayed.

Clearing UDP broadcast counters on an interface using ACLI

Use the following procedure to clear the UDP broadcast counters on an interface.

Procedure

1. Log on to ACLI in Privileged Exec command mode.
2. At the command prompt, enter the following command:

```
clear ip forward-protocol udp counters <1–4094>
```

Variable definitions

The following table describes the parameters for the **`clear ip forward-protocol udp counters`**command.

Variable	Value
<1–4094>	Specifies the VLAN ID.

Chapter 10: Directed broadcasts configuration using ACLI

This chapter describes the procedures you can use to configure and display the status of directed broadcasts using ACLI.

Configuring directed broadcasts using ACLI

Use the following procedure to enable directed broadcasts on the switch.

Note:

By default, directed broadcasts are disabled.

Before you begin

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the VLAN to be configured as a broadcast interface.
- Ensure that a route (local or static) to the destination address is available on the switch.

Procedure

1. Log on to ACLI in Global Configuration command mode.
2. At the command prompt, enter the following command:
   ```
   ip directed-broadcast enable
   ```

Displaying the directed broadcast configuration using ACLI

Use the following procedure to display the status of directed broadcasts on the switch.

Note:

By default, directed broadcasts are disabled.

Procedure

1. Log on to ACLI in User EXEC command mode.
2. At the command prompt, enter the following command:
   `show ip directed-broadcast`

Chapter 11: Static ARP and Proxy ARP configuration using ACLI

This chapter describes the procedures you can use to configure Static ARP, Proxy ARP, and display ARP entries using the ACLI.

Configuring a static ARP entry using ACLI

Use this procedure to configure a static ARP entry.

Before you begin

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the target VLAN.

Procedure

1. Log on to ACLI in Global Configuration command mode.
2. At the command prompt, enter the following command:
   ```
   [no] ip arp <A.B.C.D> <aa:bb:cc:dd:ee:ff> <unit/port> [vid <1–4094>]
   ```

Variable definitions

The following table describes the parameters for the **`ip arp`** command.

Variable	Value
[no]	Removes the specified ARP entry.
<A.B.C.D>	Specifies the IP address of the device being set as a static ARP entry.
<aa:bb:cc:dd:ee:ff>	Specifies the MAC address of the device being set as a static ARP entry.
<unit/port>	Specifies the unit and port number to which the static ARP entry is being added.

Variable	Value
vid <1–4094>	Specifies the VLAN ID to which the static ARP entry is being added.

Displaying ARP entries using ACLI

Use the following procedure to display ARP entries.

Procedure

1. Log on to ACLI in User Exec command mode.
2. At the command prompt, enter the following command:

   ```
   show arp-table
   ```

   OR

   ```
   show ip arp [<ip-addr>] [-s <subnet> <mask>] [static <ip-addr> [-s <subnet> <mask>]] [<mac-addr>] [dynamic <ip-addr> [-s <subnet> <mask>]] [<mac-addr>] {<mac_addr>] {summary] [vlan <1–4096>]
   ```

   Note:

   The `show ip arp` command is invalid if the switch is not in Layer 3 mode.

   The following information is displayed:

   - IP Address — Specifies the IP address of the ARP entry.
   - Age (min) — Displays the ARP age time.
   - MAC Address — Specifies the MAC address of the ARP entry.
   - VLAN-Unit/Port/Trunk — Specifies the VLAN/port of the ARP entry.
   - Flags — Specifies the type of ARP entry: S=Static, D=Dynamic, L=Local, B=Broadcast.

Variable definitions

The following table describes the parameters for the **`show ip arp`** command.

Variable	Value
dynamic <ip-addr> [-s <subnet> <mask>]	Displays dynamic entries for the specified subnet. If you do not specify a subnet, all dynamic entries are displayed.
<ip-addr>	Specifies the IP address of the ARP entry to be displayed.
<mac-addr>	Specifies the MAC address of the ARP entry to be displayed. The format can be H.H.H, xx:xx:xx:xx:xx:xx, xx.xx.xx.xx.xx.xx, or xx-xx-xx-xx-xx-xx-.
—s <subnet> <mask>	Displays ARP entries for the specified subnet only.
static <ip-addr> [-s <subnet> <mask>]	Displays static entries for the specified subnet. If you do not specify a subnet, all configured static entries are displayed, including those without a valid route.
summary	Displays a summary of ARP entries.
vlan <1–4096>	Displays ARP entries for a specific VLAN.

Configuring a global timeout for ARP entries using ACLI

Use the following procedure to configure an aging time for the ARP entries.

Procedure

1. Log on to ACLI in Global Configuration command mode.
2. At the command prompt, enter the following command:
   `ip arp timeout <timeout>`

Variable definitions

The following table describes the parameters for the **`ip arp timeout`** command.

Variable	Value
timeout	Specifies the amount of time in minutes before an ARP entry ages out. DEFAULT: 360 minutes. RANGE: 5–360.

Clearing the ARP cache using ACLI

Use the following procedure to clear the cache of ARP entries.

Procedure

1. Log on to ACLI in Global Configuration command mode.
2. At the command prompt, enter the following command:
   `clear arp-cache`

Configuring proxy ARP status using ACLI

Use this procedure to enable proxy ARP functionality on a VLAN.

Note:

By default, proxy ARP is disabled.

Before you begin

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the VLAN to be configured as a Proxy ARP interface.

Procedure

1. Log on to ACLI in VLAN Interface Configuration command mode.

2. At the command prompt, enter the following command:
   `[default] [no] ip arp-proxy enable`

Variable definitions

The following table describes the parameters for the **`ip arp-proxy enable`** command.

Variable	Value
[default]	Disables proxy ARP functionality on the VLAN.
[no]	Disables proxy ARP functionality on the VLAN.

Displaying proxy ARP status on a VLAN using ACLI

Use the following procedure to display the status of proxy ARP on a VLAN.

Procedure

1. Log on to ACLI in User EXEC command mode.
2. At the command prompt, enter the following command:
   `show ip arp-proxy interface [vlan <vid>]`

   The following information is displayed:

   - Vlan — Identifies a VLAN.
   - Proxy ARP status — Specifies the status of Proxy ARP on the VLAN.

Variable definitions

The following table describes the parameters for the **`show ip arp-proxy interface`** command.

Variable	Value
<vid>	Specifies the ID of the VLAN to display. RANGE:

Variable	Value
	1–4094.

Chapter 12: IGMP snooping configuration using ACLI

This chapter describes the procedures you can use to configure and display IGMP snooping parameters using ACLI.

Configuring IGMP snooping on a VLAN using ACLI

Enable IGMP snooping on a VLAN to forward the multicast data to only those ports that are members of the multicast group.

Note:

IGMP snooping is disabled by default.

Procedure

1. Log on to ACLI in VLAN Interface Configuration command mode.
2. At the command prompt, enter the following command:
   `[default] [no] ip igmp snooping`

   Note:

   Enabling IGMP snooping on a VLAN can also be done via the Global Configuration command mode. At the command prompt, enter the following command :**`[default] [no] ip igmp <vid> snooping {enable | disable}`**

Variable definitions

The following table describes the parameters for the **`ip igmp snooping`** command.

Variable	Value
default	Disables IGMP snooping on the selected VLAN.

Variable	Value
no	Disables IGMP snooping on the selected VLAN.
<vid>	Specifies the VLAN ID.
enable	Enables IGMP snooping on the selected VLAN.
disable	Disables IGMP snooping on the selected VLAN.

Enabling or disabling IGMP Multicast no flood using ACLI

Use the following procedure to enable or disable IGMP Multicast no flood.

Procedure

1. Log on to ACLI in Global Configuration command mode.
2. At the command prompt, enter the following command:
   `vlan igmp unknown-mcast-no-flood {enable|disable}`

Displaying IGMP Multicast no flood status using ACLI

Use the following procedure to display IGMP Multicast no flood status.

Procedure

1. Log on to ACLI in Global Configuration command mode.
2. At the command prompt, enter the following command:
   `show vlan igmp unknown-mcast-no-flood`

Example

```
3524GT-PWR+>enable
3524GT-PWR+#show vlan igmp unknown-mcast-no-flood
Unknown Multicast No-Flood:  Disabled
3524GT-PWR+#
```

Variable definitions

The following table describes the parameters for the **`ip igmp unknown-mcast-no-flood`** command.

Variable	Value
show	Shows the status of IGMP Multicast no flood feature.
enable	Enables IGMP Multicast no flood.
disable	Disables IGMP Multicast no flood.

Configuring IGMP proxy on a VLAN using ACLI

Use the following procedure to enable IGMP proxy on a snoop-enabled VLAN. With IGMP proxy enabled, the switch consolidates incoming report messages into one proxy report for that group.

Note:

IGMP proxy is disabled by default.

Before you begin

Enable snoop on the VLAN.

Procedure

1. Log on to ACLI in VLAN Interface Configuration command mode.
2. At the command prompt, enter the following command:
   ```
   [default] [no] ip igmp proxy
   ```

   Note:

   IGMP proxy can be enabled from Global Configuration mode by entering `[default] [no] ip igmp <vid> proxy {enable | disable}`

Variable definitions

The following table describes the parameters for the **`ip igmp proxy`** command.

Variable	Value
default	Disables IGMP proxy on the selected VLAN.
no	Disables IGMP proxy on the selected VLAN.
<vid>	Specifies the VLAN ID.
enable	Enables IGMP proxy on the selected VLAN.
disable	Disables IGMP proxy on the selected VLAN.

Configuring static mrouter ports on a VLAN using ACLI

IGMP snoop considers the port on which the IGMP query is received as the active IGMP multicast router (mrouter) port.

To forward the IGMP reports to additional ports, you can configure the additional ports as static mrouter ports.

Note:

By default, the switch forwards incoming IGMP Membership Reports only to the active mrouter port.

Procedure

1. To configure static mrouter ports on a VLAN (IGMPv1, IGMPv2, and IGMPv3 according to the supported version). Log on to ACLI in VLAN Interface Configuration command mode.
2. At the command prompt, enter the following command:
   `[default] [no] ip igmp mrouter <portlist>`

   Note:

   IGMPv1 or IGMPv2 static mrouter ports can be configured from the Global Configuration command mode by entering:`[no] ip igmp <vid> {v1-members | v2-members} {add | remove} <portlist>`

Variable definitions

The following table describes the parameters for the **`[default] [no] ip igmp mrouter`** command.

Variable	Value
default	Removes all static mrouter ports.
[no]	Removes the specified static mrouter ports.
<vid>	Specifies the VLAN on which to add the static mrouter ports.
{v1-members | v2- members}	Specifies whether the static mrouter ports are IGMPv1 or IGMPv2.
<portlist>	Specifies the list of ports to add or remove as static mrouter ports.

Configuring IGMP parameters on a VLAN using ACLI

Use the following procedure to configure the IGMP parameters on a VLAN.

Important:

The query interval and robustness values must be the same as those configured on the interface (VLAN) of the IGMP querier router.

Procedure

1. Log on to ACLI in VLAN Interface Configuration command mode.
2. At the command prompt, enter the following command:
   ```
   [default] ip igmp [last-member-query-interval <last-mbr-query-int>] [query-interval <query-int>] [query-max-response <query-max-resp>] {robust-value <robust-val>] [version <1–3>]
   ```

   Note:

   IGMP parameters can be configured from the Global Configuration command mode by entering :`[default] ip igmp <vid> [query-interval <query-int>] [robust-value <robust-val>]`

Variable definitions

The following table describes the parameters for the **`ip igmp [query-interval] [robust-value]`** command.

Variable	Value
default	Sets the selected parameter to the default value. If no parameters are specified, snoop is disabled and all IGMP parameters are set to their defaults.
<last-mbr-query-int>	Sets the maximum response time (in 1/10 seconds) that is inserted into group-specific queries sent in response to leave group messages. This parameter is also the time between group-specific query messages. This value is not configurable for IGMPv1. Decreasing the value reduces the time to detect the loss of the last member of a group. RANGE: 0–255 DEFAULT: 10 (1 second) **Note:** It is recommended to configure this parameter to values higher than 3. If a fast leave process is not required, it is recommended to have a value above 10. (The value 3 is equal to 0.3 of a second, and 10 is equal to 1.0 second.)
<query-int>	Sets the frequency (in seconds) at which host query packets are transmitted on the VLAN. RANGE: 1–65535 DEFAULT: 125 seconds
<query-max-resp>	Specifies the maximum response time (in 1/10 seconds) advertised in IGMPv2 general queries on this interface. RANGE: 0–255 DEFAULT: 100 (10 seconds)

Variable	Value
<robust-val>	Specifies tuning for the expected packet loss of a network. This value is equal to the number of expected query packet losses for each serial query interval, plus 1. **Note:** If a network is expected to lose query packets, increase the robustness value and ensure that the robustness value is equal to the configured value on the multicast router (IGMP querier). RANGE: 0–255 DEFAULT: 2 (meaning that one query for each query interval can be dropped without aging out).

Displaying IGMP interface information using ACLI

Use the following procedure to display IGMP interface information.

Procedure

1. Log on to ACLI in Privileged Exec command mode.
2. At the command prompt, enter the following command:
   `show ip igmp interface [vlan <vid>]` OR `show ip igmp <vid>`

   The following information is displayed with the `show ip igmp interface` command:

   - VLAN — Indicates the VLAN on which IGMP is configured.
   - Query Intvl — Specifies the frequency (in seconds) at which host query packets are transmitted on the interface.
   - Vers — Specifies the version of IGMP configured on the interface.
   - Oper Vers — Specifies the version of IGMP running on this interface.
   - Querier —
   - Query MaxRspT — Indicates the maximum query response time (in tenths of a second) advertised in IGMPv2 queries on this interface.
   - Wrong Query — Indicates the number of queries received whose IGMP version does not match the interface version. You must configure all routers

on a LAN to run the same version of IGMP. Thus, if queries are received with the wrong version, a configuration error occurs.

- Joins — Indicates the number of times a group membership was added on this interface.
- Robust — Specifies the robust value configured for expected packet loss on the interface.
- LastMbrQuery — Indicates the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages, and is also the amount of time between group-specific query messages. Use this value to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group. This does not apply if the interface is configured for IGMPv1.
- Send Query — Indicates whether the ip igmp send-query feature is enabled or disabled. Values are YES or NO. Default is disabled.

The following information is displayed with the `show ip igmp` command.

- Snooping — Indicates whether snooping is enabled or disabled.
- Proxy — Indicates whether proxy snoop is enabled or disabled.
- Robust Value — Indicates the robustness value configured for expected packet loss on the interface.
- Query Time — Indicates the frequency (in seconds) at which host query packets are transmitted on the interface.
- IGMPv1 Static Router Ports — Indicates the IGMPv1 static mrouter ports.
- IGMPv2 Static Router Ports — Indicates the IGMPv2 static mrouter ports.
- Send Query — Indicates whether the ip igmp send-query feature is enabled or disabled. Values are YES or NO. Default is disabled.

Variable definitions

The following table describes the parameters for the **`show ip igmp`** command.

Variable	Value
[vid <vid>]	Specifies the VLAN ID for which to display IGMP information. RANGE: 1–4094

Displaying IGMP group membership information using ACLI

Use the following procedure to display IGMP group membership information and to show the learned multicast groups and attached ports.

Procedure

1. Log on to ACLI in Privileged Exec command mode.
2. At the command prompt, enter the following command:

   `show ip igmp group [count] [group <A.B.C.D>] [member-subnet <A.B.C.D>/<0–32>]`

   `show vlan multicast membership <vid>`

   The following information is displayed after the `show ip igmp group` command:

   - Group Address — Indicates the multicast group address.
   - VLAN — Indicates the VLAN interface on which the group exists.
   - Member Address — Indicates the IP address of the IGMP receiver (host or IGMP reporter). The IP address is 0.0.0.0 if the type is static.
   - Expiration — Indicates the time left before the group report expires. This variable Is updated upon receiving a group report.
   - Type — Specifies the type of membership : static or dynamic
   - In Port — Identifies the member port for the group. This is the port on which group traffic is forwarded, and in those cases where the type is dynamic, it is the port on which the IGMP join was received.

   The following information is displayed after the `show vlan multicast membership` command:

   - Multicast Group Address — Indicates the multicast group address
   - In Port — Indicates the physical interface or the logical interface (VLAN) that received group reports from various sources.

Variable definitions

The following table describes the parameters for the **show ip igmp group** command.

Variable	Value
Group Address	Indicates the multicast group address.
VLAN	Indicates the VLAN interface on which the group exists.
Member Address	Indicates the IP address of the IGMP receiver (host or IGMP reporter). The IP address is 0.0.0.0 if the type is static.
Expiration	Indicates the time left before the group report expires. This variable is updated upon receiving a group report.
Type	Specifies the type of membership: static or dynamic
In Port	Identifies the member port for the group. This is the port on which group traffic is forwarded, and in those cases where the type is dynamic, it is the port on which the IGMP join was received

Displaying IGMP cache Information using ACLI

Use the following procedure to show the learned multicast groups in the cache and the IGMPv1 version timers.

Note:

Using the show ip igmp cache command may not display the expected results in some configurations. If the expected results are not displayed, use the show ip igmp group command to view the information.

Procedure

1. Log on to ACLI in Privileged Exec command mode.
2. At the command prompt, enter the following command:
   ```
   show ip igmp cache
   ```
   Teh following information is displayed:
   - Group Address — Indicates the multicast group address.

- VLAN ID — Indicates the VLAN interface on which the group exists.
- Last Reporter — Indicates the last IGMP host to join the group.
- Expiration — Indicates the group expiration time (in seconds).
- V1 Host Timer — Indicates the time remaining until the local router assumes that no IGMP version 1 members exist on the IP subnet attached to the interface. Upon hearing an IGMPv1 membership report, this value is reset to the group membership timer. When the time remaining is nonzero, the local interface ignores any IGMPv2 Leave messages that it receives for this group.
- Type — Indicates whether the entry is learned dynamically or is added statically.

Flushing the IGMP router table using ACLI

Use the following procedure to flush the IGMP router table.

Procedure

1. Log on to ACLI in Global Configuration command mode.
2. At the command prompt, enter the following command:
   ```
   ip igmp flush vlan <vid> {grp-member|mrouter}
   ```

Variable definitions

The following table describes the parameters for the **`ip igmp flush vlan`** command.

Variable	Value
{grp-member|mrouter}	Flushes the table specified by type.

Configuring IGMP router alert on a VLAN using ACLI

Use the following procedure to enable the router alert feature.

This feature instructs the router to drop control packets that do not have the router-alert flag in the IP header.

Note:

To maximize your network performance, it is recommended that you set the router alert option according to the version of IGMP currently in use:

- IGMPv1 — Disable
- IGMPv2 — Enable
- IGMPv3 — Enable

Procedure

1. Log on to ACLI in VLAN Interface Configuration command mode.
2. At the command prompt, enter the following command:
   `[default] [no] ip igmp router-alert`

Variable definitions

The following table describes the parameters for the **`ip igmp router-alert`** command.

Variable	Value
default	Disables the router alert option.
no	Disables the router alert option.

Chapter 13: IP routing configuration using Enterprise Device Manager

This chapter describes the procedures you can use to configure routable VLANs using Enterprise Device Manager.

The Avaya Ethernet Routing Switch 3500 Series, are Layer 3 switches. This means that a regular Layer 2 VLAN becomes a routable Layer 3 VLAN if an IP address is attached to the VLAN. When routing is enabled in Layer 3 mode, every Layer 3 VLAN is capable of routing as well as carrying the management traffic. You can use any Layer 3 VLAN instead of the Management VLAN to manage the switch.

Configuring global IP routing status and ARP lifetime using EDM

Use this procedure to enable and disable global routing at the switch level and to configure the ARP lifetime.

By default, routing is disabled.

Procedure

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IP**.
3. In the IP work area, click the **Globals** tab.
4. In the Globals section, configure Forwarding and ARPLife Time as required.
5. On the toolbar, click **Apply**.
6. On the toolbar, you can click **Refresh** verify the configuration.

Variable definitions

The following table describes the variables associated with configuring global routing and ARP lifetime.

Variable	Value
Forwarding	Indicates whether routing is enabled (forwarding) or disabled (nonforwarding) on the switch.
DefaultTTL	Indicates the default time-to-live (TTL) value for a routed packet. TTL is the maximum number of seconds elapsed before a packet is discarded. The value is inserted in the TTL field of the IP header of datagrams when one is not supplied by the transport layer protocol. The TTL field is also reduced by one each time the packet passes through a router. RANGE: 1–255 DEFAULT: 64 seconds
ReasmTimeout	Indicates the maximum number of seconds that received fragments are held while they await reassembly at this entity. DEFAULT: 60 seconds
ARPLifeTime	Specifies the lifetime in minutes of an ARP entry within the system. RANGE: 5–360 DEFAULT: 360 minutes

Configuring an IP address and enabling routing for a VLAN

Use the following procedure to configure an IP address and enable routing for a VLAN.

Procedure

1. From the navigation tree, double-click **VLAN**.

2. In the VLAN tree, click **VLANs**.
3. In the VLAN work area, select a VLAN by clicking the applicable row.
4. On the toolbar, click **IP**
5. On the toolbar, click **Insert**.
6. In the Insert IP Address section, configure as required.
7. Click **Insert**.

Variable definitions

The following table describes the variables associated with the Insert IP Address field.

Variable	Value
IpAddress	Specifies the IP address to associate with the selected VLAN.
NetMask	Specifies the subnet mask.
MacOffset	Specifies the value used to calculate the VLAN MAC address, which is offset from the switch MAC address. RANGE: 1–256 Specify the value 1 for the Management VLAN only. If no MAC offset is specified, the switch applies one automatically.

Displaying configured IP Addresses using EDM

Use the following procedure to display configured IP addresses on the switch.

Procedure

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IP**.
3. In the IP work area, click the **Addresses** tab.

Variable definitions

The following table describes the variables associated with displaying IP addresses.

Variable	Value
IfIndex	Specifies the VLAN ID.
IpAddress	Specifies the associated IP address.
NetMask	Specifies the subnet mask.
BcastAddrFormat	Specifies the format of the IP broadcast address.
ReasmMaxSize	Specifies the size of the largest IP datagram that this entity can reassemble from fragmented datagrams received on this interface.
VlanId	Specifies the VLAN ID number. A value of -1 indicates that the VLAN ID is ignored.
MacOffset	Specifies the value used to calculate the VLAN MAC address, which is offset from the switch MAC address.

Chapter 14: Static route configuration using Enterprise Device Manager

This chapter describes the procedures you can use to configure static routes using Enterprise Device Manager.

IP route management using EDM

Use the following procedures to display and filter IP route information.

Displaying IP routes using EDM

Use the following procedure to display the different routes known to the switch.

Routes are not be displayed until at least one port in the VLAN has link.

Procedure

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IP**.
3. In the IP work area, click the **Routes** tab.

Variable definitions

The following table describes the variables associated with displaying IP route information.

Variable	Value
Dest	Specifies the destination address of the route.
Mask	Specifies the subnet mask for the route.
NextHop	Specifies the next hop for the route.

Variable	Value
HopOrMetric	Specifies the metric associated with the route.
Interface	Specifies the interface associated with the route.
Proto	Specifies the protocol associated with the route. For this release, options are local or static.
PathType	Specifies the route path type: • i— indirect • d — direct • B — best • U — unresolved
Pref	Specifies the preference value associated with the route.

Filtering route information using EDM

Use the following procedure to filter the routes displayed in the Routes tab to display only the desired switch routes.

Procedure

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IP**.
3. In the IP work area, click the **Routes** tab.
4. On the toolbar, click **Filter**.
5. In the Filter route section, configure as required.
6. Click **Filter**.

Variable definitions

The following table describes the variables associated with filtering route information.

Variable	Value
Condition	When using multiple filter expressions on the tab, this is the condition that is used to join them together.
Ignore Case	Indicates whether filters are case sensitive or insensitive.
Column	Indicates the type of criteria to apply to values used for filtering.
All Records	Select this check box to clear any filters and display all rows.
Dest	Select this check box and enter a value to filter on the route destination value.
Mask	Select this check box and enter a value to filter on the route destination subnet mask value.
NextHop	Select this check box and enter a value to filter on the route next hop value.
HopOrMetric	Select this check box and enter a value to filter on the hop count or metric of the route.
Interface	Select this check box and enter a value to filter on the interface associated with the route.
Proto	Select this check box and enter a value to filter on the route protocol.
PathType	Select this check box and enter a value to filter on the route path type.
Pref	Select this check box and enter a value to filter on the route preference value.

Configuring static routes using EDM

Use the following procedure to configure static routes for the switch.

Before you begin

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the VLANs to be routed.

Procedure

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IP**.
3. In the IP work area, click the **Static Routes** tab.
4. On the toolbar, click **Insert**.
5. In the Insert Static routes section, configure as required.
6. Click **Insert**.

Variable definitions

The following table describes the variables associated with configuring static routes.

Variable	Value
Dest	Specifies the destination IP address of the route. DEFAULT: 0.0.0.0
Mask	Specifies the destination mask of the route.
NextHop	Specifies the IP address of the next hop of this route.
Metric	Represents the cost of the static route. It is used to choose the best route (the one with the smallest cost) to a certain destination. If this metric is not used, the value is set to -1. RANGE: 1–65535

Variable	Value
IfIndex	Specifies the interface on which the static route is configured.
Enable	Specifies whether the route is administratively enabled (true) or disabled (false).
Status	Specifies the operational status of the route.

Displaying TCP information for the switch using EDM

Use the following procedure to display Transmission Control Protocol (TCP) information for the switch.

Procedure

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **TCP/UDP**.
3. In the TCP/UDP work area, click the **TCP Globals** tab.

Variable definitions

The following table describes the variables associated with displaying TCP information for the switch.

Variable	Value
RtoAlgorithm	Specifies the algorithm used to determine the timeout value used for retransmitting unacknowledged octets.
RtoMin	Specifies the minimum value permitted by a TCP implementation for the retransmission timeout, measured in milliseconds.
RtoMax	Specifies the maximum value permitted by a TCP implementation for the retransmission timeout, measured in milliseconds.
MaxConn	Specifies the limit on the total number of TCP connections that the entity can support. In entities where the maximum number of

Variable	Value
	connections is dynamic, this object contains the value -1.

Displaying TCP Connections using EDM

Use the following procedure to display information about the current TCP connections.

Procedure

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **TCP/UDP**.
3. In the TCP/UDP work area, click the **TCP Connections** tab.

Variable definitions

The following table describes the variables associated with TCP connections.

Variable	Value
LocalAddressType	Specifies the local IP address type for this TCP connection.
LocalAddress	Specifies the local IP address for this TCP connection. In the case of a connection in the listen state, which is willing to accept connections for any IP interface associated with the node, the value 0.0.0.0 is used.
LocalPort	Specifies the local port number for this TCP connection.
RemAddressType	Specifies the remote IP address type for this TCP connection.
RemAddress	Specifies the remote IP address for this TCP connection.
RemPort	Specifies the remote port number for this TCP connection.
State	Specifies the state of this TCP connection.

Displaying TCP Listeners using EDM

Use the following procedure to display information about the current TCP listeners on the switch.

Procedure

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **TCP/UDP**.
3. In the TCP/UDP work area, click the **TCP Listeners** tab.

Variable definitions

The following table describes the variables associated with TCP listeners.

Variable	Value
LocalAddressType	Specifies the IP address type of the local TCP listener.
LocalAddress	Specifies the local IP address of the TCP listener. The value of this field can be represented in three possible ways, depending on the characteristics of the listening application: • For an application willing to accept both IPv4 and IPv6 datagrams, the value of this object is a zero-length octet string, and the value of the corresponding LocalAddressType field is unknown. • For an application willing to accept either IPv4 or IPv6 datagrams, the value of this object must be 0.0.0.0 or ::, with the LocalAddressType identifying the supported address type. • For an application that is listening for data destined only to a specific IP address, the value of this object is the specific local address, with LocalAddressType identifying the supported address type.

Variable	Value
LocalPort	Specifies the local port number for this TCP connection.

Displaying UDP endpoints using EDM

Use the following procedure to display information about the UDP endpoints.

Procedure

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **TCP/UDP**.
3. In the TCP/UDP work area, click the **UDP Endpoints** tab.
4. On the toolbar, you can click **Refresh** to refresh the information displayed.

Variable definitions

The following table describes the variables associated with UDP endpoints.

Variable	Value
LocalAddressType	Specifies the local address type (IPv6 or IPv4).
LocalAddress	Specifies the local IP address for this UDP listener. In the case of a UDP listener that accepts datagrams for any IP interface associated with the node, the value 0.0.0.0 is used. The value of this field can be represented in three possible ways: • For an application willing to accept both IPv4 and IPv6 datagrams, the value of this object is a zero-length octet string, and the value of the corresponding LocalAddressType field is unknown. • For an application willing to accept either IPv4 or IPv6 datagrams, the value of this object must be 0.0.0.0 or ::, with the LocalAddressType identifying the supported address type. • For an application that is listening for data destined only to a specific IP address, the

Variable	Value
	value of this object is the address for which this node is receiving packets, with LocalAddressType identifying the supported address type.
LocalPort	Specifies the local port number for this UDP listener.
RemoteAddressType	Displays the remote address type (IPv6 or IPv4).
RemoteAddress	Displays the remote IP address for this UDP endpoint. If datagrams from all remote systems are to be accepted, this value is a zero-length octet string. Otherwise, the address of the remote system from which datagrams are to be accepted (or to which all datagrams are to be sent) is displayed with the RemoteAddressType identifying the supported address type.
RemotePort	Displays the remote port number. If datagrams from all remote systems are to be accepted, this value is zero.
Instance	Distinguishes between multiple processes connected to the same UDP endpoint.
Process	Displays the ID for the UDP process.

Chapter 15: DHCP relay configuration using Enterprise Device Manager

This chapter describes the procedures you can use to configure DHCP relay using Enterprise Device Manager.

DHCP relay configuration procedures

To configure DHCP using Enterprise Device Manager, perform the following steps:

1. Specify DHCP relay configuration.
2. Specify the remote DHCP server as the destination.
3. Enable DHCP relay on the VLAN.

Enabling or disabling DHCP Forwarding using EDM

Use the following procedure to enable or disable DHCP forwarding.

Procedure

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **DHCP Relay**.
3. In the DHCP Relay work area, click the **DHCP Globals** tab.
4. In the DhcpForwardingEnabled section, check box to enable or uncheck box to disable.
5. On the toolbar, click **Apply**.

Configuring DHCP Forwarding maximum frame size globally using EDM

Use the following procedure to specify the maximum frame size the DHCP relay agent can forward to the DHCP server.

Procedure

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **DHCP Relay**.
3. In the DHCP Relay work area, click the **DHCP Globals** tab.
4. In the DhcpForwardingMaxFrameLength section, enter the frame length between 576 and 1536 bytes.

   Note:

   The default value is 576 bytes.
5. On the toolbar, click **Apply**.

Configuring DHCP Relay using EDM

Use this procedure to configure DHCP Relay.

Procedure

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **DHCP Relay**.
3. In the DHCP Relay work area, click the **DHCP Relay** tab.
4. Click **Insert**.
5. In the Insert section, configure as required.
6. Click **Insert**.
7. On the toolbar, you can click **Refresh** to verify the configuration.

Variable definitions

The following table describes the variables associated with configuring the DHCP relay.

Variable	Value
AgentAddr	The IP address of the local VLAN serving as the DHCP relay agent.
ServerAddr	The IP address of the remote DHCP server.
Enable	Enables (selected) or disables (cleared) DHCP relay.
Mode	Indicates whether the relay instance applies for BOOTP packets, DHCP packets, or both.

Configuring DHCP Relay with Option 82 globally using EDM

Use this procedure to enable DHCP Relay Option 82 globally.

Procedure

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **DHCP Relay**.
3. In the DHCP Relay work area, click the **DHCP Globals** tab.
4. In the DhcpForwardingOption82Enabled section, check the box to enable.
5. On the toolbar, click **Apply**.

Configuring DHCP Relay with Option 82 for a VLAN using EDM

Use this procedure to configure DHCP Relay with Option 82 for a VLAN.

Before you begin

- Enable IP routing globally.
- On the VLAN: enable IP Routing and configure an IP address to be set as the DHCP Relay agent.
- Ensure that a route, either local or static, is available on the switch to the destination DHCP server.

Procedure

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **DHCP Relay**.
3. In the DHCP Relay work area, click the **DHCP Relay-VLAN** tab.
4. In the table, double click the **Option82Enabled** cell to edit.
    - **true** enables DHCP Relay with Option 82 for the VLAN
    - **false** disables DHCP Relay with Option 82 for the VLAN
5. On the toolbar, click **Apply**.

Configuring DHCP parameters on a VLAN using EDM

Use the following procedure to configure the DHCP relay parameters on a VLAN.

Procedure

1. From the navigation tree, double-click **VLAN**.
2. In the VLAN tree, click **VLANs**.
3. In the VLANs work area, click the **Basic** tab.
4. In the Basic section, select the VLAN for which the DHCP relay is to be configured.
5. On the toolbar, click **IP**.

6. Select the **DHCP** tab.
7. In the DHCP section, configure as required.
8. Click **Apply**.

Variable definitions

The following table describes the variables associated with DHCP parameters on VLANs.

Variable	Value
Enable	Specifies whether DHCP relay is enabled or disabled.
MinSec	Indicates the min-sec value. The switch immediately forwards a BootP/DHCP packet if the secs field in the BootP/DHCP packet header is greater than the configured min-sec value; otherwise, the packet is dropped.
Mode	Specifies the type of packets this VLAN interface forwards: BootP, DHCP, or both.
AlwaysBroadcast	Specifies whether DHCP Reply packets are broadcast to the DHCP clients on this VLAN interface.
ClearCounters	Specifies to clear the DHCP relay counters for the VLAN.
CounterClearTime	Specifies the last time the counter values in this entry were reset to 0.

Configuring DHCP Relay with Option 82 for a VLAN using EDM

Use this procedure to configure DHCP Relay with Option 82 for a VLAN.

Before you begin

- Enable IP routing globally.
- On the VLAN: enable IP Routing and configure an IP address to be set as the DHCP Relay agent.
- Ensure that a route, either local or static, is available on the switch to the destination DHCP server.

Procedure

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **DHCP Relay**.
3. In the DHCP Relay work area, click the **DHCP Relay-VLAN** tab.
4. In the table, double click the **Option82Enabled** cell to edit.
    - **true** enables DHCP Relay with Option 82 for the VLAN
    - **false** disables DHCP Relay with Option 82 for the VLAN
5. On the toolbar, click **Apply**.

Displaying and graphing DHCP counters on a VLAN using EDM

Use the following procedure to display and graph the current DHCP counters on a VLAN.

Procedure

1. From the navigation tree, double-click **VLAN**.
2. In the VLAN tree, click **VLANs**.
3. In the VLANs work area, click the **Basic** tab.
4. In the Basic section, select a VLAN.
5. On the toolbar, click **IP**.

6. In the **IP** work area, click the **DHCP** tab.
7. Click **Graph**.
8. On the toolbar, select a **Poll interval** from the drop down menu.
9. Select **Line, Area, Bar** or **Pie** chart.
   The following information is displayed:
   - NumRequests — indicates the number of DHCP requests.
   - NumReplies — indicates the number of DHCP replies.

Assigning a DHCP Relay Option 82 subscriber ID to a port using EDM

Use the following procedure to assign a DHCP Relay Option 82 subscriber ID to a port.

Before you begin

- Enable IP Routing globally.
- On the VLAN: enable IP Routing and configure an IP address to be set as the DHCP Relay agent.
- Ensure the a route, either local or static, is available on the switch to the destination DHCP server.

Procedure

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **DHCP Relay**.
3. In the DHCP Relay work area, click the **DHCP Relay-port** tab.
4. In the Multiple Port Configuration section, click the ellipsis and highlight required port(s), click **OK**.
5. In the PortDhcpOption82SubscriberId section, double click cell and enter **subscriber ID** for the port.
6. On the toolbar, click **Apply**.

Variable definitions

The following table describes the variables associated with Option 82 subscriber ID.

Variable	Value
rcPortIndex	Indicates the slot and port number.
PortDhcpOption82SubscriberId	Specifies the DHCP Option 82 subscriber ID for the port. The value is a character string between 1 and 64.

Chapter 16: UDP broadcast forwarding configuration using Enterprise Device Manager

This chapter describes the procedures you can use to configure and manage UDP broadcast forwarding using Enterprise Device Manager.

UDP broadcast forwarding is a general mechanism for selectively forwarding limited UDP broadcasts received on an IP interface to a configured IP address.

UDP broadcast forwarding configuration procedures

To configure UDP broadcast forwarding using Enterprise Device Manager, perform the following steps:

1. Create UDP protocol entries that specify each UDP port and associated protocol that you want to forward.
2. Create UDP forwarding entries that specify the destination address for each UDP port that you want to forward.
3. Add UDP forwarding entries to a UDP forwarding list (you can create up to 128 UDP forwarding lists.)
4. Apply UDP forwarding lists to local VLAN interfaces.

Configuring UDP protocol table entries using EDM

Use the following procedure to create UDP table entries that identify the protocols associated with specific UDP ports to forward.

Procedure

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **UDP Forwarding**.
3. In the UDP Forwarding area, click the **Protocols** tab.

4. In the Protocols section, click **Insert**.
5. In the Insert Protocols section, configure as required.
6. Click **Insert**.

Variable definitions

The following table describes the variables associated with configuring UDP protocol table entries.

Variable	Value
PortNumber	Specifies the UDP port number.
Name	Specifies the protocol name associated with the UDP port.

Configuring UDP forwarding entries using EDM

Use the following procedure to configure individual UDP forwarding entries, which associate UDP forwarding ports with destination IP addresses.

Procedure

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **UDP Forwarding**.
3. In the UDP Forwarding work area, click the **Forwardings** tab.
4. On the toolbar, click **Insert**.
5. In the Insert Forwardings section, specify a destination address.
6. Click **Insert**.

Variable definitions

The following table describes the variables associated with UDP forward entries.

Name	Description
DestPort	Specifies the port on which the UDP forwarding originates (configured using the Protocols tab).
DestAddress	Specifies the destination IP address.

Configuring a UDP forwarding list using EDM

Use the following procedure to add the UDP port/destination forwarding entries (configured in the Forwardings tab) to UDP forwarding lists.

Each UDP forwarding list can contain multiple port/destination entries.

Procedure

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **UDP Forwarding**.
3. In the UDP Forwarding work area, click the **Forwarding Lists** tab.
4. On the toolbar, click **Insert**.
5. In the Insert Fowarding Lists section, configure as required.
6. In the **FwdIdList** section, click the ellipsis and select the desired port/destination pairs.
7. Click **Ok**.
8. Click **Insert**.

Variable definitions

The following table describes the variables associated with UDP forwarding lists.

Variable	Value
Id	The unique identifier assigned to the forwarding list.
Name	The name assigned to the forwarding list.
FwdIdList	The forwarding entry IDs associated with the port/server IP pairs created using the Forwardings tab.

Applying a UDP forwarding list to a VLAN using EDM

Use the following procedure to assign a UDP forwarding list to a VLAN and to configure the related UDP forwarding parameters for the VLAN.

Procedure

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **UDP Forwarding**.
3. In the UDP Forwarding work area, click the **Broadcast Interfaces** tab.
4. Click **Insert**.
5. In the Insert Broadcast Interface section, configure as required.
6. Click **Insert**.

Variable definitions

The following table describes the variables associated with applying a UDP forwarding list to a VLAN.

Variable	Value
LocalIfAddr	Specifies the IP address of the local VLAN interface.

Variable	Value
UdpPortFwdListId	Specifies the port forwarding lists associated with the interface. This ID is defined in the Forwarding Lists tab.
MaxTtl	Indicates the maximum number of hops an IP broadcast packet can take from the source device to the destination device. This is an integer value between 1 and 16.
NumRxPkts	Specifies the total number of UDP broadcast packets received by this local interface.
NumFwdPkts	Specifies the total number of UDP broadcast packets forwarded.
NumDropPktsDestUnreach	Specifies the total number of UDP broadcast packets dropped because the destination is unreachable.
NumDropPktsUnknownPort	Specifies the total number of UDP broadcast packets dropped because the destination port or protocol specified has no matching forwarding policy.
BroadCastMask	Specifies the 32-bit mask used by the selected VLAN interface to take forwarding decisions based on the destination IP address of the incoming UDP broadcast traffic. If you do not specify a broadcast mask value, the switch uses the mask of the interface to which the forwarding list is attached.

Chapter 17: Static ARP and Proxy ARP configuration using Enterprise Device Manager

This chapter describes the procedures you can use to configure Static ARP, display ARP entries, and configure Proxy ARP using Enterprise Device Manager.

Configuring static ARP entries using EDM

Use this procedure to configure static ARP entries for the switch.

Before you begin

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the target VLAN interface.

Procedure

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IP**.
3. In the IP work area, click the **ARP** tab.
4. Click **Insert**.
5. Click **Port in Vlan** and select the VLAN to add the static ARP entry.
6. Configure entries as required.
7. Click **Insert**.

Variable definitions

The following table describes the variables associated with configuring static ARP entries.

Variable	Value
Interface	Specifies the VLAN and port to which the static ARP entry is being added.
MacAddress	Specifies the MAC address of the device being set as a static ARP entry.
IpAddress	Specifies the IP address of the device being set as a static ARP entry.
Type	Specifies the type of ARP entry: static, dynamic, or local.

Configuring Proxy ARP using EDM

Use the following procedure to configure proxy ARP on the switch. Proxy ARP allows the switch to respond to an ARP request from a locally attached host (or end station) for a remote destination.

Procedure

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IP**.
3. In the IP work area, click the **ARP Interfaces** tab.

   Important:

   Device Manager does not display the ARP Interfaces tab if you have not enabled routing on the switch.
4. In the ARP Interfaces section, click the **DoProxy column** on a VLAN.
5. Click **Enable**.
6. Click **Apply**.

Variable definitions

The following table describes the variables associated with the ARP interface tab.

Variable	Value
IfIndex	Specifies a configured switch interface.
DoProxy	Enables or disables proxy ARP on the interface.
DoResp	Specifies whether the sending of ARP responses on the specified interface is enabled or disabled.

Chapter 18: IGMP snooping configuration using Enterprise Device Manager

This chapter describes the procedures used to configure IGMP snooping using Enterprise Device Manager.

Managing IGMP snoop using EDM

Use the following procedures to configure IGMP snooping and proxy and static mrouter ports.

Configuring IGMP snooping using EDM

Use the following procedure to configure IGMP snooping on a switch.

Procedure

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IGMP**.
3. In the IGMP work area, click the **Snoop** tab.
4. In the work area, configure as required.
   To add static mrouter ports, specify the desired ports as follows:
   - **Ver1MRouterPorts** field (for IGMP version 1)
   - **Ver2MRouterPorts** field (for IGMP version 2)
   - **MRouterPorts** field (for both IGMP versions)
5. On the toolbar, click **Apply**.

Variable definitions

The following table describes the variables associated with IGMP snooping.

Variable	Value
Id	Specifies the VLAN ID.
Name	Specifies the VLAN name.
Enable	Specifies whether IGMP snooping is enabled or disabled.
ReportProxyEnable	Specifies whether IGMP proxy is enabled or disabled.
Robustness	Specifies tuning for the expected packet loss of a network. This value is equal to the number of expected query packet losses for each serial query interval, plus 1. If you expect a network to lose query packets, you must increase the robustness value. Ensure that the robustness value is the same as the configured value on the multicast router (IGMP querier). RANGE: 0–255 DEFAULT: 2 The default value of 2 means that one query for each query interval can be dropped without the querier aging out.
QueryInterval	Sets the frequency (in seconds) at which host query packets are transmitted on the VLAN.
MRouterPorts	Specifies ports in the VLAN that provide connectivity to an IP Multicast router.
Ver1MRouterPorts	Specifies ports in this VLAN that provide connectivity to an IP Multicast router using IGMP version 1.
Ver2MRouterPorts	Specifies ports in this VLAN that provide connectivity to an IP Multicast router using IGMP version 2.
ActiveMRouterPorts	Specifies the active mrouter ports (dynamic and static) in this VLAN that provide connectivity to an IP Multicast router.
ActiveQuerier	Specifies the IP address of the multicast querier router.

Variable	Value
QuerierPort	Specifies the port on which the multicast querier router is heard.
MRouterExpiration	Specifies the multicast querier router aging timeout.

Configuring IGMP snoop proxy and static mrouter ports on a VLAN using EDM

Use the following procedure to configure IGMP snooping, proxy, and static mrouter ports on a VLAN.

By default, IGMP snoop and proxy are disabled, and no static mrouter ports are configured.

Procedure

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IGMP**.
3. In the IGMP work area, click the **Snoop** tab.
4. In the Snoop section, configure cells as required.
5. On the toolbar, click **Apply**.
6. On the toolbar, you can click **Refresh** verify the configuration.

Variable definitions

The following table describes the fields on the IGMP snoop tab.

Variable	Value
IfIndex	Specifies the VLAN ID.
SnoopEnable	Specifies the IGMP snoop status: • enabled (true) • disabled (false)
ProxySnoopEnable	Specifies the IGMP proxy status: • enabled (true) • disabled (false)

Variable	Value
SnoopMRouterPorts	Specifies the static mrouter ports. Such ports are directly attached to a multicast router so the multicast data and group reports are forwarded to the router.
SnoopActiveMRouterPorts	Displays all dynamic (querier port) and static mrouter ports that are active on the interface.
SnoopMRouterExpiration	Specifies the time remaining before the multicast router is aged out on this interface. If the switch does not receive queries before this time expires, it flushes out all group memberships known to the VLAN. The Query Max Response Interval (obtained from the queries received) is used as the timer resolution.

Displaying IGMP groups using EDM

Use this procedure to display the IGMP group information.

Procedure

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IGMP**.
3. In the IGMP work area, click the **Groups** tab.

Variable definitions

The following table describes the variables associated with IGMP group information.

Variable	Value
IpAddress	Indicates the multicast group IP address. An address can be the same for many incoming ports.
IfIndex	Indicates VLAN interface associated with the multicast group address.
Members	Indicates the IP address of the IGMP receiver (host or IGMP reporter).

Variable	Value
Expiration	Indicates the time left before the group report expires. This variable is updated when a group report is received.
InPort	Indicates the member port for the group. This is the port on which group traffic is forwarded.

Displaying IGMP group information using EDM

Use the following procedure to display IGMP group information.

Procedure

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IGMP**.
3. In the IGMP work area, click the **Groups-Ext** tab.

Variable definitions

The following table describes the variables associated with IGMP group information.

Variable	Value
IpAddress	Indicates the multicast group address.
SourceAddress	Indicates the source address.
Members	Indicates the IP address of the IGMP receiver (host or IGMP reporter).
Mode	Indicates the mode.
IfIndex	Indicates the VLAN interface from which the multicast group address is heard.
Expiration	Indicates the time left before the group report expires on this port. This variable is updated upon receiving a group report.
InPort	Indicates the member port for the group. This is the port on which group traffic is forwarded.

Displaying IGMP cache information using EDM

Use the following procedure to display IGMP cache information.

Procedure

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IGMP**.
3. In the IGMP work area, click the **Cache** tab.

Variable definitions

The following table describes the variables associated with IGMP cache information.

Name	Description
Address	Indicates the IP multicast group address.
IfIndex	Indicates the VLAN interface from which the group address is heard.
LastReporter	Indicates the last IGMP host to join the group.
ExpiryTime	Indicates the amount of time (in seconds) remaining before this entry is aged out.
Version1Host Timer	Indicates the time remaining until the local router assumes that no IGMP version 1 members exist on the IP subnet attached to the interface. Upon hearing an IGMPv1 membership report, this value is reset to the group membership timer. When the time remaining is nonzero, the local interface ignores IGMPv2 Leave messages that it receives for this group.
Type	Indicates whether the entry is learned dynamically or is added statically.

Managing IP Address multicast filter tables using EDM

Use the following procedures to display IP address multicast filter tables and specify IP address flooding.

Specifying an IP address to be allowed to flood a VLAN using EDM

Use the following procedure to configure the IP address multicast filter table. This table specifies multicast IP addresses that are allowed to be flooded to all ports on a per-VLAN basis.

Procedure

1. From the navigation tree, double-click **VLAN**.
2. In the VLAN tree, click **VLANs**.
3. In the VLANs work area, click the **IP address Multicast Filter Table** tab.
4. Click **Insert**.
5. In the Insert section, configure as required.
6. Click **Insert**.

Variable definitions

The following table describes the variables of the IP Address Multicast Filter Table tab.

Variable	Value
VlanAllowedInetAddressVlanId	Specifies the ID of the VLAN to configure.
VlanAllowedInetAddressType	Specifies the address type: ipv4.
VlanAllowedInetAddress	Specifies a multicast IP address that is allowed to flood all ports. Unicast and broadcast addresses are not allowed.

Displaying the IP Address Multicast Filter Table using EDM

Use the following procedure to display the IP Multicast Filter Table.

Procedure

1. From the navigation tree, double-click **VLAN**.
2. In the VLAN tree, click **VLANs**.
3. In the VLANs work area, click the **IP Address Multicast Filter Table** tab.

Variable definitions

The following table describes the variables associated with the IP Address Multicast Filter Table.

Variable	Value
VlanAllowedInetAddressVlanId	The ID of the VLAN in which the specified multicast IP address is allowed to flood traffic.
VlanAllowedInetAddressVlanType	The address type. The only supported value is ipv4.
VlanAllowedInetAddress	Multicast IP address. Traffic destined to this address will be flooded inside the VLAN.

Configuring IGMP interface parameters and flushing IGMP tables using EDM

Use the following procedure to make interface specific IGMP settings and/or flush the IGMP tables on a VLAN.

Procedure

1. From the navigation tree, double-click **IP**.
2. In the IP tree, click **IGMP**.
3. In the IGMP work area, click the **Interface** tab.
4. Double click the cell under the **FlushAction** column and select the desired flush option.
5. On the toolbar, click **Apply**.

Variable definitions

The following table describes the fields on the IGMP Interface tab.

Name	Description
IfIndex	Indicated the interface on which the IGMP is enabled.
QueryInterval	Indicates the frequency (in seconds) at which IGMP host query packets are transmitted on the interface. Ensure that the robustness value is the same as the configured value on the multicast router (IGMP querier). RANGE: 1–65535 DEFAULT: 125
Status	Indicates whether or not the interface is active. The interface becomes active if any IGMP forwarding ports exist on the interface. If the VLAN has no port members or if all of the port members are disabled, the status is notInService.
Version	Indicates the version of IGMP (1, 2, or 3) configured on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. DEFAULT: 2
OperVersion	Indicates the version of IGMP currently running on this interface.
Querier	Indicates the address of the IGMP querier on the IP subnet to which this interface is attached.
QueryMaxResponseTime	Indicates the maximum response time (in 1/10 seconds) advertised in IGMPv2 general queries on this interface.
WrongVersionQueries	Indicates the number of queries received with an IGMP version that does not match the interface. IGMP requires that all routers on a LAN be configured to run the same version of IGMP. If queries are received with the wrong version, it indicates a version mismatch.
Joins	Indicates the number of times a group membership is added on this interface; that is, the number of times an entry for this

Name	Description
	interface is added to the cache table. This number gives an indication of the amount of IGMP activity over time.
Robustness	Specifies tuning for the expected packet loss of a network. This value is equal to the number of expected query packet losses for each serial query interval, plus 1. If you expect a network to lose query packets, you must increase the robustness value. Ensure that the robustness value is the same as the configured value on the multicast router (IGMP querier). RANGE: 2–255 DEFAULT: 2 The default value of 2 means that one query for each query interval can be dropped without the querier aging out.
LastMembQueryIntvl	Sets the maximum response time (in tenths of a second) that is inserted into group-specific queries sent in response to leave group messages. This parameter is also the time between groupspecific query messages. This value is not configurable for IGMPv1. Decreasing the value reduces the time to detect the loss of the last member of a group. RANGE: 0–255 Avaya recommends configuring this parameter to values higher than 3. If a fast leave process is not required, Avaya recommends values above 10. (The value 3 is equal to 0.3 of a second, and 10 is equal to 1.0 second.)
RouterAlertEnable	When enabled, this parameter instructs the router to ignore IGMP packets that do not contain the router alert IP option. When disabled (default setting), the router processes IGMP packets regardless of whether the router alert IP option is set or not. To maximize your network performance, Avaya recommends thatyou set this parameter according to the version of IGMP currently in use:

Name	Description
	• IGMPv1—Disable • IGMPv2—Enable • IGMPv3—Enable
SendQuery	Indicates whether to enable the SendQuery feature on this vlan or not. With SendQuery enabled, a multicast snooping capable switch will send out general queries at every query interval, overcoming the absence of an actual mrouter in the LAN.
FlushAction	Flushes the specified table type: • none • flushGrpMem — group member table • flushMrouter — mrouter table

Enabling or disabling IGMP Multicast no flood using EDM

Use the following procedure to enable or disable IGMP Multicast no flood.

Procedure

1. From the navigation tree, double-click **VLAN**.
2. In the VLAN tree, click **VLANs**.
3. In the VLANs work area, click the **Unknown Multicast Filtering** tab.
4. Select the **UnknownMulticastNoFlood** check box to enable or clear the check box to disable.
5. Click **Apply**.

Enabling or disabling IGMP Multicast no flood using ACLI

Use the following procedure to enable or disable IGMP Multicast no flood.

Procedure

1. Log on to ACLI in Global Configuration command mode.
2. At the command prompt, enter the following command:

```
vlan igmp unknown-mcast-no-flood {enable|disable}
```

Configuring VLAN snooping using EDM

Use this procedure to configure VLAN snooping.

Procedure

1. From the navigation tree, double-click **VLAN**.
2. In the VLAN tree, click **VLANs**.
3. In the VLANs work area, click the **Snoop** tab.
4. In the Snoop section, configure as required.
5. On the toolbar, click **Apply**.

Variable definitions

The following table describes the fields on the VLAN snoop tab.

Name	Description
Id	Specifies the VLAN ID.
Name	Specifies the VLAN name.
Enable	Specifies whether snooping is enabled or disabled.
ReportProxyEnable	Specifies whether the proxy is enabled or disabled.
Robustness	Specifies tuning for the expected packet loss of a network. This value is equal to the number of expected query packet losses for each serial query interval, plus 1. If you expect a network to lose query packets, you must increase the robustness value. Ensure that the robustness value is the same as the configured value on the multicast router. RANGE: 0–255 DEFAULT: 2

Name	Description
	The default value of 2 means that one query for each query interval can be dropped without the querier aging out.
QueryInterval	Sets the frequency (in seconds) at which host query packets are transmitted on the VLAN.
MRouterPorts	Specifies ports in the VLAN that provide connectivity to an IP Multicast router.
Ver1MRouterPorts	Specifies ports in this VLAN that provide connectivity to an IP Multicast router using IGMP version 1.
Ver2MRouterPorts	Specifies ports in this VLAN that provide connectivity to an IP Multicast router using IGMP version 2.
ActiveMRouterPorts	Specifies the active mrouter ports (dynamic and static) in this VLAN that provide connectivity to an IP Multicast router.
ActiveQuerier	Specifies the IP address of the multicast querier router.
QuerierPort	Specifies the port on which the multicast querier router is heard.
MRouterExpiration	Specifies the multicast querier router aging timeout.

Displaying the MAC Multicast Filter Table using EDM

Use the following procedure to display the MAC Multicast Filter Table.

Procedure

1. From the navigation tree, double-click **VLAN**.
2. In the VLAN tree, click **VLANs**.
3. In the VLANs work area, click the **MAC Multicast Filter Table** tab.

Variable definitions

The following table describes the variables associated with the Multicast Filter Table.

Name	Description
AllowedAddressMacAddr	Indicates the MAC addresses for which flooding is allowed.
AllowedAddressVlanId	Indicates the VLAN interface for which the multicast MAC address is allowed.