


User Manual



EKI-9728G Series

Ind. Rackmount L3 Managed
Switch with AC/DC

ADVANTECH

Enabling an Intelligent Planet

Copyright

The documentation and the software included with this product are copyrighted 2016 by Advantech Co., Ltd. All rights are reserved. Advantech Co., Ltd. reserves the right to make improvements in the products described in this manual at any time without notice. No part of this manual may be reproduced, copied, translated or transmitted in any form or by any means without the prior written permission of Advantech Co., Ltd. Information provided in this manual is intended to be accurate and reliable. However, Advantech Co., Ltd. assumes no responsibility for its use, nor for any infringements of the rights of third parties, which may result from its use.

Acknowledgements

Intel and Pentium are trademarks of Intel Corporation.

Microsoft Windows and MS-DOS are registered trademarks of Microsoft Corp.

All other product names or trademarks are properties of their respective owners.

Product Warranty (5 years)

Advantech warrants to you, the original purchaser, that each of its products will be free from defects in materials and workmanship for five years from the date of purchase.

This warranty does not apply to any products which have been repaired or altered by persons other than repair personnel authorized by Advantech, or which have been subject to misuse, abuse, accident or improper installation. Advantech assumes no liability under the terms of this warranty as a consequence of such events.

Because of Advantech's high quality-control standards and rigorous testing, most of our customers never need to use our repair service. If an Advantech product is defective, it will be repaired or replaced at no charge during the warranty period. For out-of-warranty repairs, you will be billed according to the cost of replacement materials, service time and freight. Please consult your dealer for more details.

If you think you have a defective product, follow these steps:

1. Collect all the information about the problem encountered. (For example, CPU speed, Advantech products used, other hardware and software used, etc.) Note anything abnormal and list any on screen messages you get when the problem occurs.
2. Call your dealer and describe the problem. Please have your manual, product, and any helpful information readily available.
3. If your product is diagnosed as defective, obtain an RMA (return merchandise authorization) number from your dealer. This allows us to process your return more quickly.
4. Carefully pack the defective product, a fully-completed Repair and Replacement Order Card and a photocopy proof of purchase date (such as your sales receipt) in a shippable container. A product returned without proof of the purchase date is not eligible for warranty service.
5. Write the RMA number visibly on the outside of the package and ship it prepaid to your dealer.

Part No. XXXXXXXXXXXX

Printed in Taiwan

Edition 1

December 2017

Declaration of Conformity

CE

This product has passed the CE test for environmental specifications when shielded cables are used for external wiring. We recommend the use of shielded cables. This kind of cable is available from Advantech. Please contact your local supplier for ordering information.

This product has passed the CE test for environmental specifications. Test conditions for passing included the equipment being operated within an industrial enclosure. In order to protect the product from being damaged by ESD (Electrostatic Discharge) and EMI leakage, we strongly recommend the use of CE-compliant industrial enclosure products.

FCC Class A

Note: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his own expense.

FCC Class A

Note: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his own expense.

FM

This equipment has passed the FM certification. According to the National Fire Protection Association, work sites are classified into different classes, divisions and groups, based on hazard considerations. This equipment is compliant with the specifications of Class I, Division 2, Groups A, B, C and D indoor hazards.

Technical Support and Assistance

1. Visit the Advantech web site at www.advantech.com/support where you can find the latest information about the product.
2. Contact your distributor, sales representative, or Advantech's customer service center for technical support if you need additional assistance. Please have the following information ready before you call:
 - Product name and serial number
 - Description of your peripheral attachments
 - Description of your software (operating system, version, application software, etc.)
 - A complete description of the problem
 - The exact wording of any error messages

Warnings, Cautions and Notes

Warning! *Warnings indicate conditions, which if not observed, can cause personal injury!*



Caution! *Cautions are included to help you avoid damaging hardware or losing data. e.g.*



There is a danger of a new battery exploding if it is incorrectly installed. Do not attempt to recharge, force open, or heat the battery. Replace the battery only with the same or equivalent type recommended by the manufacturer. Discard used batteries according to the manufacturer's instructions.

Note! *Notes provide optional additional information.*



Document Feedback

To assist us in making improvements to this manual, we would welcome comments and constructive criticism. Please send all such - in writing to: support@advantech.com

Packing List

Before setting up the system, check that the items listed below are included and in good condition. If any item does not accord with the table, please contact your dealer immediately.

- 1 x Full Managed Ethernet Switch
- 3 x Terminal Blocks
- 1 x Startup Manual

Safety Instructions

1. Read these safety instructions carefully.
2. Keep this User Manual for later reference.
3. Disconnect this equipment from any AC outlet before cleaning. Use a damp cloth. Do not use liquid or spray detergents for cleaning.
4. For plug-in equipment, the power outlet socket must be located near the equipment and must be easily accessible.
5. Keep this equipment away from humidity.
6. Put this equipment on a reliable surface during installation. Dropping it or letting it fall may cause damage.
7. The openings on the enclosure are for air convection. Protect the equipment from overheating. **DO NOT COVER THE OPENINGS.**
8. Make sure the voltage of the power source is correct before connecting the equipment to the power outlet.
9. Position the power cord so that people cannot step on it. Do not place anything over the power cord.
10. All cautions and warnings on the equipment should be noted.
11. If the equipment is not used for a long time, disconnect it from the power source to avoid damage by transient overvoltage.
12. Never pour any liquid into an opening. This may cause fire or electrical shock.
13. Never open the equipment. For safety reasons, the equipment should be opened only by qualified service personnel.
14. If one of the following situations arises, get the equipment checked by service personnel:
 15. The power cord or plug is damaged.
 16. Liquid has penetrated into the equipment.
 17. The equipment has been exposed to moisture.
 18. The equipment does not work well, or you cannot get it to work according to the user's manual.
 19. The equipment has been dropped and damaged.
 20. The equipment has obvious signs of breakage.
21. **DO NOT LEAVE THIS EQUIPMENT IN AN ENVIRONMENT WHERE THE STORAGE TEMPERATURE MAY GO BELOW -20° C (-4° F) OR ABOVE 60° C (140° F). THIS COULD DAMAGE THE EQUIPMENT. THE EQUIPMENT SHOULD BE IN A CONTROLLED ENVIRONMENT.**
22. **CAUTION: DANGER OF EXPLOSION IF BATTERY IS INCORRECTLY REPLACED. REPLACE ONLY WITH THE SAME OR EQUIVALENT TYPE RECOMMENDED BY THE MANUFACTURER, DISCARD USED BATTERIES ACCORDING TO THE MANUFACTURER'S INSTRUCTIONS.**
23. The sound pressure level at the operator's position according to IEC 704-1:1982 is no more than 70 dB (A).

DISCLAIMER: This set of instructions is given according to IEC 704-1. Advantech disclaims all responsibility for the accuracy of any statements contained herein.

Wichtige Sicherheitshinweise

1. Bitte lesen sie Sich diese Hinweise sorgfältig durch.
2. Heben Sie diese Anleitung für den späteren Gebrauch auf.
3. Vor jedem Reinigen ist das Gerät vom Stromnetz zu trennen. Verwenden Sie Keine Flüssig-oder Aerosolreiniger. Am besten dient ein angefeuchtetes Tuch zur Reinigung.
4. Die Netzanschlussteckdose soll nahe dem Gerät angebracht und leicht zugänglich sein.
5. Das Gerät ist vor Feuchtigkeit zu schützen.
6. Bei der Aufstellung des Gerätes ist auf sicheren Stand zu achten. Ein Kippen oder Fallen könnte Verletzungen hervorrufen.
7. Die Belüftungsöffnungen dienen zur Luftzirkulation die das Gerät vor überhitzung schützt. Sorgen Sie dafür, daß diese Öffnungen nicht abgedeckt werden.
8. Beachten Sie beim. Anschluß an das Stromnetz die Anschlußwerte.
9. Verlegen Sie die Netzanschlusbleitung so, daß niemand darüber fallen kann. Es sollte auch nichts auf der Leitung abgestellt werden.
10. Alle Hinweise und Warnungen die sich am Geräten befinden sind zu beachten.
11. Wird das Gerät über einen längeren Zeitraum nicht benutzt, sollten Sie es vom Stromnetz trennen. Somit wird im Falle einer Überspannung eine Beschädigung vermieden.
12. Durch die Lüftungsöffnungen dürfen niemals Gegenstände oder Flüssigkeiten in das Gerät gelangen. Dies könnte einen Brand bzw. elektrischen Schlag auslösen.
13. Öffnen Sie niemals das Gerät. Das Gerät darf aus Gründen der elektrischen Sicherheit nur von autorisiertem Servicepersonal geöffnet werden.
14. Wenn folgende Situationen auftreten ist das Gerät vom Stromnetz zu trennen und von einer qualifizierten Servicestelle zu überprüfen:
 15. Netzkabel oder Netzstecker sind beschädigt.
 16. Flüssigkeit ist in das Gerät eingedrungen.
 17. Das Gerät war Feuchtigkeit ausgesetzt.
 18. Wenn das Gerät nicht der Bedienungsanleitung entsprechend funktioniert oder Sie mit Hilfe dieser Anleitung keine Verbesserung erzielen.
 19. Das Gerät ist gefallen und/oder das Gehäuse ist beschädigt.
 20. Wenn das Gerät deutliche Anzeichen eines Defektes aufweist.
21. **VORSICHT:** Explosionsgefahr bei unsachgemäßen Austausch der Batterie.Ersatz nur durch denselben oder einem vom Hersteller empfohlene-männlichen Typ. Entsorgung gebrauchter Batterien nach Angaben des Herstellers.
22. **ACHTUNG:** Es besteht die Explosionsgefahr, falls die Batterie auf nicht fachmännische Weise gewechselt wird. Verfangen Sie die Batterie nur gleicher oder entsprechender Type, wie vom Hersteller empfohlen. Entsorgen Sie Batterien nach Anweisung des Herstellers.
23. Der arbeitsplatzbezogene Schalldruckpegel nach DIN 45 635 Teil 1000 beträgt 70dB(A) oder weiger.

Haftungsausschluss: Die Bedienungsanleitungen wurden entsprechend der IEC-704-1 erstellt. Advantech lehnt jegliche Verantwortung für die Richtigkeit der in diesem Zusammenhang getätigten Aussagen ab.

Safety Precaution - Static Electricity

Follow these simple precautions to protect yourself from harm and the products from damage.

- To avoid electrical shock, always disconnect the power from your PC chassis before you work on it. Don't touch any components on the CPU card or other cards while the PC is on.
- Disconnect power before making any configuration changes. The sudden rush of power as you connect a jumper or install a card may damage sensitive electronic components.

Contents

Chapter 1	Product Overview	1
1.1	Specifications	2
1.2	Hardware Views	3
1.2.1	Front View	3
1.2.2	Rear View	3
1.2.3	Dimensions	5
Chapter 2	Switch Installation	6
2.1	Warnings	7
2.2	Installation Guidelines	9
2.3	Environment and Enclosure Guidelines	9
2.3.1	Connecting Hardware	9
2.4	Verifying Switch Operation	10
2.5	Installing the Switch	10
2.5.1	Rack-Mounting	10
2.6	Installing and Removing SFP Modules	11
2.6.1	Installing SFP Modules	11
2.6.2	Removing SFP Modules	13
2.7	Connecting the Switch to Ethernet Ports	14
2.7.1	RJ45 Ethernet Cable Wiring	14
2.8	Connecting the Switch to Console Port	14
2.9	Power Supply Installation	15
2.9.1	Overview	15
2.9.2	Considerations	16
2.9.3	Grounding the Device	16
2.9.4	Wiring a Relay Contact	17
2.9.5	Wiring the Power Inputs	17
2.10	Reset Button	18
Chapter 3	Configuration Utility	19
3.1	First Time Setup	20
3.1.1	Overview	20
3.1.2	Introduction	20
3.1.3	Administrative Interface Access	20
3.1.4	Using the Graphical (Web) Interface	21
3.1.5	Configuring the Switch for Network Access	21
3.1.6	Configuring the Ethernet Ports	22
3.2	Command Line Interface Configuration	23
3.2.1	Introduction to Command-Line Interface (CLI)	23
3.2.2	Accessing the CLI	23
3.3	Web Browser Configuration	24
3.3.1	Preparing for Web Configuration	24
3.3.2	System Login	24
Chapter 4	Managing Switch	25
4.1	Log In	26
4.2	Recommended Practices	26
4.2.1	Changing Default Password	26

4.3	System	27
	4.3.1 AAA.....	27
	4.3.2 Advanced Configuration.....	35
	4.3.3 Basic Configuration.....	97
	4.3.4 Configuration Storage.....	98
	4.3.5 Connectivity	100
	4.3.6 Firmware.....	107
	4.3.7 Logs.....	110
	4.3.8 Management Access	117
	4.3.9 Management Security.....	124
	4.3.10 Passwords	125
	4.3.11 Port	128
	4.3.12 Slot.....	134
	4.3.13 Statistics.....	136
	4.3.14 Status.....	147
	4.3.15 Summary.....	150
	4.3.16 Users.....	154
	4.3.17 Utilities	160
4.4	Switching.....	173
	4.4.1 Auto Recovery	173
	4.4.2 Class of Service	175
	4.4.3 DHCP Snooping.....	176
	4.4.4 IPv6 DHCP Snooping	188
	4.4.5 DVLAN.....	194
	4.4.6 Dynamic ARP Inspection	196
	4.4.7 Filters	202
	4.4.8 GARP.....	204
	4.4.9 IGMP Snooping.....	206
	4.4.10 IGMP Snooping Querier.....	212
	4.4.11 MLD Snooping.....	215
	4.4.12 MLD Snooping Querier	221
	4.4.13 Multicast Forwarding Database	224
	4.4.14 MVR	228
	4.4.15 LLDP.....	232
	4.4.16 LLDP-MED.....	237
	4.4.17 Port Channel.....	241
	4.4.18 Port Security	243
	4.4.19 Protected Ports	249
	4.4.20 Spanning Tree	250
	4.4.21 UDLD	261
	4.4.22 VLAN.....	263
	4.4.23 IP Subnet Based VLAN.....	269
	4.4.24 MAC Based VLAN	270
	4.4.25 Protocol Based VLAN	272
	4.4.26 Private VLAN	275
	4.4.27 Voice VLAN.....	278
	4.4.28 Virtual Port Channel.....	280
4.5	Routing.....	286
	4.5.1 ARP Table.....	286
	4.5.2 IP.....	289
	4.5.3 Router	297
	4.5.4 IPv6.....	302
	4.5.5 IPv6 Routes	316
	4.5.6 DHCPv6.....	320
4.6	Security	327
	4.6.1 Port Access Control	327
	4.6.2 RADIUS	337
	4.6.3 TACACS+	343
	4.6.4 Authentication Manager.....	346

4.7	QoS.....	350
4.7.1	Access Control Lists	350
4.7.2	Auto VoIP.....	366
4.7.3	Class of Service.....	370
4.7.4	Diffserv.....	373
Chapter 5	Command Line Interface	387
5.1	Virtual Router Redundancy Protocol Commands	388
5.1.1	Virtual Router Redundancy Protocol Commands	388
5.2	Open Shortest Path First Commands	394
5.2.1	General OSPF Commands.....	394
5.2.2	OSPF Interface Commands.....	410
5.2.3	IP Event Dampening Commands	415
5.2.4	OSPF Graceful Restart Commands	416
5.2.5	OSPFv2 Stub Router Commands.....	420
5.2.6	OSPF Show Commands.....	421
5.3	Routing Information Protocol Commands	440
5.3.1	Routing Information Protocol Commands.....	440
Chapter 6	Troubleshooting.....	447
6.1	Troubleshooting	448

List of Figures

Figure 1.1	Front View	3
Figure 1.2	Front View	3
Figure 1.3	System LED Panel	4
Figure 1.4	Dimensions	5
Figure 2.1	Installing the Rack Mount Brackets	10
Figure 2.2	Installing the Switch	10
Figure 2.3	Removing the Dust Plug from an SFP Slot	11
Figure 2.4	Installing an SFP Transceiver	12
Figure 2.5	Attaching a Fiber Optic Cable to a Transceiver	12
Figure 2.6	Removing a Fiber Optic Cable to a Transceiver	13
Figure 2.7	Removing an SFP Transceiver	13
Figure 2.8	Ethernet Plug & Connector Pin Position	14
Figure 2.9	Serial Console Cable	14
Figure 2.10	DB 9 Pin Position	14
Figure 2.11	Pin Assignment	15
Figure 2.12	Power Wiring for EKI-9228G Series	16
Figure 2.13	Terminal Receptor: Relay Contact	17
Figure 2.14	Terminal Receptor: Power Input Contacts	18
Figure 2.15	Installing DC Wires in a Terminal Block	18
Figure 4.1	Login Screen	26
Figure 4.2	System > Users > Accounts	26
Figure 4.3	Changing a Default Password	27
Figure 4.4	System > AAA > Authentication List	27
Figure 4.5	System > AAA > Authentication List > Add	29
Figure 4.6	System > AAA > Authentication Selection	29
Figure 4.7	System > AAA > Authorization List	30
Figure 4.8	System > AAA > Authorization List > Add	31
Figure 4.9	System > AAA > Authorization Selection	32
Figure 4.10	System > AAA > Accounting List	33
Figure 4.11	System > AAA > Accounting List > Add	34
Figure 4.12	System > AAA > Accounting Selection	34
Figure 4.13	System > Advanced Configuration > DHCP Server > Global	35
Figure 4.14	System > Advanced Configuration > DHCP Server > Excluded Addresses	36
Figure 4.15	System > Advanced Configuration > DHCP Server > Excluded Addresses > Add	36
Figure 4.16	System > Advanced Configuration > DHCP Server > Pool Summary	37
Figure 4.17	System > Advanced Configuration > DHCP Server > Pool Summary > Add	38
Figure 4.18	System > Advanced Configuration > DHCP Server > Pool Configuration	39
Figure 4.19	System > Advanced Configuration > DHCP Server > Pool Options	41
Figure 4.20	System > Advanced Configuration > DHCP Server > Pool Options > Add Vendor Option	42
Figure 4.21	System > Advanced Configuration > DHCP Server > Pool Options > Edit	43
Figure 4.22	System > Advanced Configuration > DHCP Server > Bindings	43
Figure 4.23	System > Advanced Configuration > DHCP Server > Statistics	44
Figure 4.24	System > Advanced Configuration > DHCP Server > Conflicts	45
Figure 4.25	System > Advanced Configuration > DNS > Configuration	46
Figure 4.26	System > Advanced Configuration > DNS > IP Mapping	47
Figure 4.27	System > Advanced Configuration > DNS > IP Mapping > Add	48
Figure 4.28	System > Advanced Configuration > DNS > Source Interface Configuration	48
Figure 4.29	System > Advanced Configuration > Email Alerts > Global	49
Figure 4.30	System > Advanced Configuration > Email Alerts > Test	50
Figure 4.31	System > Advanced Configuration > Email Alerts > Server	51
Figure 4.32	System > Advanced Configuration > Email Alerts > Server > Add	51
Figure 4.33	System > Advanced Configuration > Email Alerts > Statistics	52
Figure 4.34	System > Advanced Configuration > Email Alerts > Subject	52
Figure 4.35	System > Advanced Configuration > Email Alerts > Address	53
Figure 4.36	System > Advanced Configuration > Email Alerts > Address > Add	53
Figure 4.37	System > Advanced Configuration > ISDP > Global	54
Figure 4.38	System > Advanced Configuration > ISDP > Cache Table	54
Figure 4.39	System > Advanced Configuration > ISDP > Interface	55
Figure 4.40	System > Advanced Configuration > ISDP > Statistics	56
Figure 4.41	System > Advanced Configuration > Link Dependency > Group	57
Figure 4.42	System > Advanced Configuration > Link Dependency > Group > Add	58
Figure 4.43	System > Advanced Configuration > LLPF > Configuration	58
Figure 4.44	System > Advanced Configuration > Protection > Denial of Service	59
Figure 4.45	System > Advanced Configuration > SDM > SDM	61
Figure 4.46	System > Advanced Configuration > sFlow > Agent	61
Figure 4.47	System > Advanced Configuration > sFlow > Receiver	62

Figure 4.48	System > Advanced Configuration > sFlow > Poller	63
Figure 4.49	System > Advanced Configuration > sFlow > Poller > Add.....	63
Figure 4.50	System > Advanced Configuration > sFlow > Sampler	64
Figure 4.51	System > Advanced Configuration > sFlow > Sampler > Add	65
Figure 4.52	System > Advanced Configuration > sFlow > Source Interface Configuration.....	65
Figure 4.53	System > Advanced Configuration > SNMP > Community	66
Figure 4.54	System > Advanced Configuration > SNMP > Community > Add Community	67
Figure 4.55	System > Advanced Configuration > SNMP > Community > Add Community Group.....	67
Figure 4.56	System > Advanced Configuration > SNMP > Trap Receiver v1/v2	68
Figure 4.57	System > Advanced Configuration > SNMP > Trap Receiver v1/v2 > Add.....	69
Figure 4.58	System > Advanced Configuration > SNMP > Trap Receiver v3	70
Figure 4.59	System > Advanced Configuration > SNMP > Trap Receiver v3 > Add	71
Figure 4.60	System > Advanced Configuration > SNMP > Supported MIBs.....	72
Figure 4.61	System > Advanced Configuration > SNMP > Access Control Group	72
Figure 4.62	System > Advanced Configuration > SNMP > Access Control Group > Add.....	73
Figure 4.63	System > Advanced Configuration > SNMP > User Security Model.....	74
Figure 4.64	System > Advanced Configuration > SNMP > User Security Model > Add	75
Figure 4.65	System > Advanced Configuration > SNMP > Source Interface Configuration.....	76
Figure 4.66	System > Advanced Configuration > SNMP > Server Configuration	77
Figure 4.67	System > Advanced Configuration > SNTP > Global Configuration	78
Figure 4.68	System > Advanced Configuration > SNTP > Global Status.....	79
Figure 4.69	System > Advanced Configuration > SNTP > Server Configuration	80
Figure 4.70	System > Advanced Configuration > SNTP > Server Configuration > Add.....	81
Figure 4.71	System > Advanced Configuration > SNTP > Server Status	81
Figure 4.72	System > Advanced Configuration > SNTP > Source Interface Configuration	82
Figure 4.73	System > Advanced Configuration > Time Ranges > Configuration	83
Figure 4.74	System > Advanced Configuration > Time Ranges > Configuration > Add	84
Figure 4.75	System > Advanced Configuration > Time Ranges > Entry Configuration.....	85
Figure 4.76	System > Advanced Configuration > Time Ranges > Entry Configuration > Add Absolute	85
Figure 4.77	System > Advanced Configuration > Time Ranges > Entry Configuration > Add Periodic	86
Figure 4.78	System > Advanced Configuration > Time Zone > Summary	87
Figure 4.79	System > Advanced Configuration > Time Zone > Time Zone	88
Figure 4.80	System > Advanced Configuration > Time Zone > Summer Time	89
Figure 4.81	System > Advanced Configuration > Trap Manager > Trap Log.....	91
Figure 4.82	System > Advanced Configuration > Trap Manager > Trap Flags	92
Figure 4.83	System > Advanced Configuration > CPU Traffic Filter > Global	92
Figure 4.84	System > Advanced Configuration > CPU Traffic Filter > Filter Configuration	93
Figure 4.85	System > Advanced Configuration > CPU Traffic Filter > Interfaces	95
Figure 4.86	System > Advanced Configuration > CPU Traffic Filter > Interfaces > Add.....	95
Figure 4.87	System > Advanced Configuration > CPU Traffic Filter > Statistics	96
Figure 4.88	System > Advanced Configuration > CPU Traffic Filter > Summary	96
Figure 4.89	System > Advanced Configuration > CPU Traffic Filter > Trace Information	97
Figure 4.90	System > Basic Configuration > Switch	97
Figure 4.91	System > Configuration Storage > Save	98
Figure 4.92	System > Configuration Storage > Reset.....	98
Figure 4.93	System > Configuration Storage > Erase Startup	99
Figure 4.94	System > Configuration Storage > Copy.....	99
Figure 4.95	System > Connectivity > IPv4	100
Figure 4.96	System > Connectivity > IPv6	101
Figure 4.97	System > Connectivity > IPv6 Neighbors	102
Figure 4.98	System > Connectivity > IPv6 Neighbors > Add.....	103
Figure 4.99	System > Connectivity > Service Port IPv4	104
Figure 4.100	System > Connectivity > Service Port IPv6	105
Figure 4.101	System > Connectivity > Service Port IPv6 Neighbors	106
Figure 4.102	System > Connectivity > Service Port IPv6 Neighbors List > Add	107
Figure 4.103	System > Connectivity > DHCP Client Options.....	107
Figure 4.104	System > Firmware > Status	107
Figure 4.105	System > Firmware > Configuration and Upgrade	108
Figure 4.106	System > Firmware > AutoInstall	109
Figure 4.107	System > Logs > Buffered Log.....	110
Figure 4.108	System > Logs > Event Log	111
Figure 4.109	System > Logs > Persistent Log	112
Figure 4.110	System > Logs > Hosts	112
Figure 4.111	System > Logs > Hosts > Add.....	113
Figure 4.112	System > Logs > Configuration	114
Figure 4.113	System > Logs > Source Interface Configuration.....	115
Figure 4.114	System > Logs > Statistics	116
Figure 4.115	System > Management Access > System.....	117
Figure 4.116	System > Management Access > Telnet.....	118

Figure 4.117	System > Management Access > Outbound Telnet	118
Figure 4.118	System > Management Access > Serial.....	119
Figure 4.119	System > Management Access > CLI Banner.....	120
Figure 4.120	System > Management Access > HTTP	120
Figure 4.121	System > Management Access > HTTPS	121
Figure 4.122	System > Management Access > SSH	122
Figure 4.123	System > Management Security > Access Profile.....	124
Figure 4.124	System > Passwords > Line Password	125
Figure 4.125	System > Passwords > Enable Password.....	126
Figure 4.126	System > Passwords > Password Rules.....	126
Figure 4.127	System > Passwords > Last Password	127
Figure 4.128	System > Passwords > Reset Passwords.....	128
Figure 4.129	System > Port > Summary	128
Figure 4.130	System > Port > Description.....	130
Figure 4.131	System > Port > Cable Test	130
Figure 4.132	System > Port > Mirroring.....	131
Figure 4.133	System > Port > Mirroring Summary	133
Figure 4.134	System > Slot > Configuration.....	134
Figure 4.135	System > Slot > Configuration > Add	134
Figure 4.136	System > Slot > Supported Cards.....	135
Figure 4.137	System > Statistics > System > Switch	136
Figure 4.138	System > Statistics > System > Port Summary.....	137
Figure 4.139	System > Statistics > System > Port Detailed	138
Figure 4.140	System > Statistics > System > Network DHCPv6.....	141
Figure 4.141	System > Statistics > Time Based > Group.....	142
Figure 4.142	System > Statistics > Time Based > Group > Add	144
Figure 4.143	System > Statistics > Time Based > Flow Based.....	145
Figure 4.144	System > Statistics > Time Based > Flow Based > Add	146
Figure 4.145	System > Statistics > Time Based > Statistics	147
Figure 4.146	System > Status > ARP Cache	147
Figure 4.147	System > Status > Resource Status.....	148
Figure 4.148	System > Status > Resource Configuration	149
Figure 4.149	System > Summary > Dashboard	150
Figure 4.150	System > Summary > Description	151
Figure 4.151	System > Summary > Inventory	152
Figure 4.152	System > Summary > MAC Address Table.....	153
Figure 4.153	System > Users > Accounts	154
Figure 4.154	System > Users > Accounts > Add.....	155
Figure 4.155	System > Users > Auth Server Users.....	156
Figure 4.156	System > Users > Auth Server Users > Add	157
Figure 4.157	System > Users > Sessions	157
Figure 4.158	System > Users > User Domain Name	158
Figure 4.159	System > Users > Task Groups	158
Figure 4.160	System > Users > Task Groups > Add	159
Figure 4.161	System > Users > User Groups.....	159
Figure 4.162	System > Users > User Groups > Add	160
Figure 4.163	System > Utilities > System Reset	160
Figure 4.164	System > Utilities > Ping.....	161
Figure 4.165	System > Utilities > Ping IPv6	162
Figure 4.166	System > Utilities > TraceRoute	163
Figure 4.167	System > Utilities > TraceRoute IPv6.....	165
Figure 4.168	System > Utilities > IP Address Conflict	166
Figure 4.169	System > Utilities > Transfer	167
Figure 4.170	System > Utilities > Digital Signature Verification.....	170
Figure 4.171	System > Utilities > Core Dump	171
Figure 4.172	System > Utilities > Core Dump Test	172
Figure 4.173	Switching > Auto Recovery > Configuration	173
Figure 4.174	Switching > Class of Service > 802.1p	175
Figure 4.175	Switching > DHCP Snooping > Base > Global.....	176
Figure 4.176	Switching > DHCP Snooping > Base > VLAN Configuration.....	176
Figure 4.177	Switching > DHCP Snooping > Base > VLAN Configuration > Add	177
Figure 4.178	Switching > DHCP Snooping > Base > Interface Configuration	177
Figure 4.179	Switching > DHCP Snooping > Base > Static Bindings.....	178
Figure 4.180	Switching > DHCP Snooping > Base > Static Bindings > Add	179
Figure 4.181	Switching > DHCP Snooping > Base > Dynamic Bindings.....	180
Figure 4.182	Switching > DHCP Snooping > Base > Persistent	180
Figure 4.183	Switching > DHCP Snooping > Base > Statistics	181
Figure 4.184	Switching > DHCP Snooping > L2 Relay > Global	182
Figure 4.185	Switching > DHCP Snooping > L2 Relay > Interface Configuration	182
Figure 4.186	Switching > DHCP Snooping > L2 Relay > VLAN Configuration.....	183
Figure 4.187	Switching > DHCP Snooping > L2 Relay > VLAN Configuration > Add	184
Figure 4.188	Switching > DHCP Snooping > L2 Relay > Statistics.....	185

Figure 4.189	Switching > DHCP Snooping > IP Source Guard > Interface Configuration	186
Figure 4.190	Switching > DHCP Snooping > IP Source Guard > Bindings.....	186
Figure 4.191	Switching > DHCP Snooping > IP Source Guard > Bindings > Add	187
Figure 4.192	Switching > IPv6 DHCP Snooping > Base > Global.....	188
Figure 4.193	Switching > IPv6 DHCP Snooping > Base > VLAN Configuration	188
Figure 4.194	Switching > IPv6 DHCP Snooping > Base > VLAN Configuration > Add.....	189
Figure 4.195	Switching > IPv6 DHCP Snooping > Base > Interface Configuration.....	190
Figure 4.196	Switching > IPv6 DHCP Snooping > Base > Static Bindings	191
Figure 4.197	Switching > IPv6 DHCP Snooping > Base > Static Bindings > Add.....	191
Figure 4.198	Switching > IPv6 DHCP Snooping > Base > Dynamic Bindings	192
Figure 4.199	Switching > IPv6 DHCP Snooping > Base > Persistent.....	192
Figure 4.200	Switching > IPv6 DHCP Snooping > Base > Statistics.....	193
Figure 4.201	Switching > DVLAN > Configuration	194
Figure 4.202	Switching > DVLAN > Summary	195
Figure 4.203	Switching > DVLAN > Interface Summary	195
Figure 4.204	Switching > Dynamic ARP Inspection > Global.....	196
Figure 4.205	Switching > Dynamic ARP Inspection > VLAN.....	197
Figure 4.206	Switching > Dynamic ARP Inspection > VLAN > Add	198
Figure 4.207	Switching > Dynamic ARP Inspection > Interface	198
Figure 4.208	Switching > Dynamic ARP Inspection > ACL Summary.....	199
Figure 4.209	Switching > Dynamic ARP Inspection > ACL Summary > Add	200
Figure 4.210	Switching > Dynamic ARP Inspection > ACL Configuration.....	200
Figure 4.211	Switching > Dynamic ARP Inspection > ACL Configuration > Add Rule.....	201
Figure 4.212	Switching > Dynamic ARP Inspection > Statistics.....	201
Figure 4.213	Switching > Filters > MAC Filters	202
Figure 4.214	Switching > Filters > MAC Filters > Add.....	203
Figure 4.215	Switching > GARP > Switch	204
Figure 4.216	Switching > GARP > Port.....	205
Figure 4.217	Switching > IGMP Snooping > Configuration	206
Figure 4.218	Switching > IGMP Snooping > Interface Configuration	207
Figure 4.219	Switching > IGMP Snooping > Source Specific Multicast	208
Figure 4.220	Switching > IGMP Snooping > VLAN Status	208
Figure 4.221	Switching > IGMP Snooping > VLAN Status > Add	209
Figure 4.222	Switching > IGMP Snooping > Multicast Router Configuration	210
Figure 4.223	Switching > IGMP Snooping > Multicast Router VLAN Status.....	211
Figure 4.224	Switching > IGMP Snooping > Multicast Router VLAN Configuration.....	211
Figure 4.225	Switching > IGMP Snooping Querier > Configuration	212
Figure 4.226	Switching > IGMP Snooping Querier > VLAN Configuration.....	213
Figure 4.227	Switching > IGMP Snooping Querier > VLAN Configuration > Add	214
Figure 4.228	Switching > IGMP Snooping Querier > VLAN Status	214
Figure 4.229	Switching > MLD Snooping > Configuration.....	215
Figure 4.230	Switching > MLD Snooping > Interface Configuration.....	216
Figure 4.231	Switching > MLD Snooping > Source Specific Multicast.....	217
Figure 4.232	Switching > MLD Snooping > VLAN Status	217
Figure 4.233	Switching > MLD Snooping > VLAN Status > Add	218
Figure 4.234	Switching > MLD Snooping > Multicast Router Configuration.....	219
Figure 4.235	Switching > MLD Snooping > Multicast Router VLAN Status	220
Figure 4.236	Switching > MLD Snooping > Multicast Router VLAN Status > Add.....	220
Figure 4.237	Switching > MLD Snooping Querier > Configuration.....	221
Figure 4.238	Switching > MLD Snooping Querier > VLAN Configuration	222
Figure 4.239	Switching > MLD Snooping Querier > VLAN Configuration > Add.....	223
Figure 4.240	Switching > MLD Snooping Querier > VLAN Status	223
Figure 4.241	Switching > Multicast Forwarding Database > Summary.....	224
Figure 4.242	Switching > Multicast Forwarding Database > GMRP	225
Figure 4.243	Switching > Multicast Forwarding Database > IGMP Snooping.....	226
Figure 4.244	Switching > Multicast Forwarding Database > Source Specific Multicast	226
Figure 4.245	Switching > Multicast Forwarding Database > Source Specific Multicast Status	227
Figure 4.246	Switching > Multicast Forwarding Database > Statistics	228
Figure 4.247	Switching > MVR > Global	228
Figure 4.248	Switching > MVR > Group.....	229
Figure 4.249	Switching > MVR > Group > Add	230
Figure 4.250	Switching > MVR > Interface.....	230
Figure 4.251	Switching > MVR > Statistics	231
Figure 4.252	Switching > LLDP > Global	232
Figure 4.253	Switching > LLDP > Interface.....	233
Figure 4.254	Switching > LLDP > Interface > Add	234
Figure 4.255	Switching > LLDP > Local Devices	235
Figure 4.256	Switching > LLDP > Remote Devices	235
Figure 4.257	Switching > LLDP > Statistics	236
Figure 4.258	Switching > LLDP-MED > Global	237
Figure 4.259	Switching > LLDP-MED > Interface.....	238
Figure 4.260	Switching > LLDP-MED > Interface > Add	238

Figure 4.261	Switching > LLDP-MED > Local Devices	239
Figure 4.262	Switching > LLDP-MED > Remote Devices	240
Figure 4.263	Switching > Port Channel > Summary	241
Figure 4.264	Switching > Port Channel > Statistics	242
Figure 4.265	Switching > Port Security > Global	243
Figure 4.266	Switching > Port Security > Interface	244
Figure 4.267	Switching > Port Security > VLAN	245
Figure 4.268	Switching > Port Security > VLAN	246
Figure 4.269	Switching > Port Security > Static MAC	247
Figure 4.270	Switching > Port Security > Static MAC > Add	247
Figure 4.271	Switching > Port Security > Dynamic MAC	248
Figure 4.272	Switching > Protected Ports > Configuration	249
Figure 4.273	Switching > Protected Ports > Configuration > Add	249
Figure 4.274	Switching > Spanning Tree > Switch	250
Figure 4.275	Switching > Spanning Tree > MST	251
Figure 4.276	Switching > Spanning Tree > MST > Add	252
Figure 4.277	Switching > Spanning Tree > MST Port	253
Figure 4.278	Switching > Spanning Tree > CST	254
Figure 4.279	Switching > Spanning Tree > CST Port	256
Figure 4.280	Switching > Spanning Tree > Statistics	257
Figure 4.281	Switching > Spanning Tree > PVST Global	258
Figure 4.282	Switching > Spanning Tree > PVST VLAN	258
Figure 4.283	Switching > Spanning Tree > PVST VLAN > Add	259
Figure 4.284	Switching > Spanning Tree > PVST Interface	259
Figure 4.285	Switching > Spanning Tree > PVST Statistics	260
Figure 4.286	Switching > UDLD > Configuration	261
Figure 4.287	Switching > UDLD > Interface Configuration	261
Figure 4.288	Switching > VLAN > Status	263
Figure 4.289	Switching > VLAN > Status > Add	264
Figure 4.290	Switching > VLAN > Port Configuration	264
Figure 4.291	Switching > VLAN > Port Summary	265
Figure 4.292	Switching > VLAN > Switchport Summary	267
Figure 4.293	Switching > VLAN > Internal Usage	268
Figure 4.294	Switching > VLAN > Reset	268
Figure 4.295	Switching > VLAN > Status	269
Figure 4.296	Switching > IP Subnet Based VLAN > Status	269
Figure 4.297	Switching > IP Subnet Based VLAN > Status > Add	270
Figure 4.298	Switching > MAC Based VLAN > Status	270
Figure 4.299	Switching > MAC Based VLAN > Status > Add	271
Figure 4.300	Switching > Protocol Based VLAN > Status	272
Figure 4.301	Switching > Protocol Based VLAN > Status > Add	273
Figure 4.302	Switching > Protocol Based VLAN > Configuration	274
Figure 4.303	Switching > Private VLAN > Configuration	275
Figure 4.304	Switching > Private VLAN > Configuration > Add VLAN	276
Figure 4.305	Switching > Private VLAN > Association	276
Figure 4.306	Switching > Private VLAN > Interface	277
Figure 4.307	Switching > Voice VLAN > Configuration	278
Figure 4.308	Switching > Voice VLAN > Interface Summary	278
Figure 4.309	Switching > Voice VLAN > Interface Summary > Add	279
Figure 4.310	Switching > Virtual Port Channel > Global	280
Figure 4.311	Switching > Virtual Port Channel > Interface Configuration	283
Figure 4.312	Switching > Virtual Port Channel > Interface Configuration > Add	284
Figure 4.313	Switching > Virtual Port Channel > Statistics	284
Figure 4.314	Routing > ARP Table > Summary	286
Figure 4.315	Routing > ARP Table > Summary > Add	287
Figure 4.316	Routing > ARP Table > Configuration	288
Figure 4.317	Routing > ARP Table > Statistics	288
Figure 4.318	Routing > IP > Configuration	289
Figure 4.319	Routing > IP > Interface Summary	290
Figure 4.320	Routing > IP > Interface Configuration	291
Figure 4.321	Routing > IP > Statistics	294
Figure 4.322	Routing > IP > Statistics	295
Figure 4.323	Routing > Router > Route Table	297
Figure 4.324	Routing > Router > Configured Routes	298
Figure 4.325	Routing > Router > Configured Routes > Add	299
Figure 4.326	Routing > Router > Summary	300
Figure 4.327	Routing > Router > ECMP Group	301
Figure 4.328	Routing > IPv6 > Configuration	302
Figure 4.329	Routing > IPv6 > Interface Summary	303
Figure 4.330	Routing > IPv6 > Interface Configuration	304
Figure 4.331	Routing > IPv6 > Loopback Configuration	306
Figure 4.332	Routing > IPv6 > Global Addresses	307

Figure 4.333	Routing > IPv6 > Address Configuration	308
Figure 4.334	Routing > IPv6 > Statistics	309
Figure 4.335	Routing > IPv6 > Detailed Statistics	310
Figure 4.336	Routing > IPv6 > Neighbor Table	314
Figure 4.337	Routing > IPv6 > Neighbor Table > Add	315
Figure 4.338	Routing > IPv6 Routes > IPv6 Route Table	316
Figure 4.339	Routing > IPv6 Routes > IPv6 Configured Routes	316
Figure 4.340	Routing > IPv6 Routes > IPv6 Configured Routes > Add	317
Figure 4.341	Routing > IPv6 Routes > IPv6 ECMP Group.....	318
Figure 4.342	Routing > IPv6 Routes > IPv6 Route Summary	319
Figure 4.343	Routing > DHCPv6 > Global	320
Figure 4.344	Routing > DHCPv6 > Pool Summary	321
Figure 4.345	Routing > DHCPv6 > Pool Summary > Add	321
Figure 4.346	Routing > DHCPv6 > Pool Configuration	322
Figure 4.347	Routing > DHCPv6 > Interface.....	323
Figure 4.348	Routing > DHCPv6 > Interface Configuration	324
Figure 4.349	Routing > DHCPv6 > Bindings	325
Figure 4.350	Routing > DHCPv6 > Statistics	326
Figure 4.351	Security > Port Access Control > Configuration	327
Figure 4.352	Security > Port Access Control > Port Summary	328
Figure 4.353	Security > Port Access Control > Port Configuration	330
Figure 4.354	Security > Port Access Control > Port Details.....	332
Figure 4.355	Security > Port Access Control > Statistics	334
Figure 4.356	Security > Port Access Control > Client Summary	335
Figure 4.357	Security > Port Access Control > Privileges Summary	335
Figure 4.358	Security > Port Access Control > History Log Summary	336
Figure 4.359	Security > RADIUS > Configuration	337
Figure 4.360	Security > RADIUS > Named Server	338
Figure 4.361	Security > RADIUS > Named Server > Add	339
Figure 4.362	Security > RADIUS > Statistics	339
Figure 4.363	Security > RADIUS > Accounting Server	340
Figure 4.364	Security > RADIUS > Accounting Server > Add.....	341
Figure 4.365	Security > RADIUS > Accounting Statistics	341
Figure 4.366	Security > RADIUS > Clear Statistics.....	342
Figure 4.367	Security > RADIUS > Source Interface Configuration	342
Figure 4.368	Security > TACACS+ > Configuration	343
Figure 4.369	Security > TACACS+ > Server Summary.....	344
Figure 4.370	Security > TACACS+ > Server Summary > Add	344
Figure 4.371	Security > TACACS+ > Server Configuration.....	345
Figure 4.372	Security > TACACS+ > Source Interface Configuration.....	345
Figure 4.373	Security > Authentication Manager > Configuration	346
Figure 4.374	Security > Authentication Manager > Authentication Tiering	347
Figure 4.375	Security > Authentication Manager > Authenticated Clients	348
Figure 4.376	Security > Authentication Manager > Statistics.....	349
Figure 4.377	Security > Authentication Manager > History	349
Figure 4.378	QoS > Access Control Lists > Summary	350
Figure 4.379	QoS > Access Control Lists > Summary > Add.....	351
Figure 4.380	QoS > Access Control Lists > Configuration	352
Figure 4.381	QoS > Access Control Lists > Configuration > Add Rule	354
Figure 4.382	QoS > Access Control Lists > Configuration > Add Rule	359
Figure 4.383	QoS > Access Control Lists > Interfaces.....	359
Figure 4.384	QoS > Access Control Lists > Interfaces > Add	360
Figure 4.385	QoS > Access Control Lists > VLANs	361
Figure 4.386	QoS > Access Control Lists > VLANs > Add.....	362
Figure 4.387	QoS > Access Control Lists > Control Plane.....	363
Figure 4.388	QoS > Access Control Lists > Control Plane > Add	364
Figure 4.389	QoS > Access Control Lists > Statistics	364
Figure 4.390	QoS > Auto VoIP > Global	366
Figure 4.391	QoS > Auto VoIP > OUI Table	366
Figure 4.392	QoS > Auto VoIP > OUI Table > Add	367
Figure 4.393	QoS > Auto VoIP > OUI Based Auto VoIP	367
Figure 4.394	QoS > Auto VoIP > Protocol Based Auto VoIP	368
Figure 4.395	QoS > Class of Service > IP DSCP.....	370
Figure 4.396	QoS > Class of Service > Interface	371
Figure 4.397	QoS > Class of Service > Queue	372
Figure 4.398	QoS > Diffserv > Global	373
Figure 4.399	QoS > Diffserv > Class Summary	374
Figure 4.400	QoS > Diffserv > Class Summary > Add.....	374
Figure 4.401	QoS > Diffserv > Class Configuration	375
Figure 4.402	QoS > Diffserv > Class Configuration > Add Match Criteria.....	376
Figure 4.403	QoS > Diffserv > Policy Summary	379
Figure 4.404	QoS > Diffserv > Policy Summary > Add	379

Figure 4.405	QoS > Diffserv > Policy Configuration	380
Figure 4.406	QoS > Diffserv > Policy Configuration > Add Class.....	380
Figure 4.407	QoS > Diffserv > Policy Configuration > Add Attribute	381
Figure 4.408	QoS > Diffserv > Service Summary.....	384
Figure 4.409	QoS > Diffserv > Service Summary > Add	385
Figure 4.410	QoS > Diffserv > Service Statistics.....	385
Figure 4.411	QoS > Diffserv > Policy Statistics	386

Chapter 1

Product Overview

1.1 Specifications

Specifications	Description		
Interface	I/O Port	24 x RJ45 + 4 x 10GbE SFP+ 8 x RJ45/SFP (mini-GBIC) combo port	
	Power Connector	<ul style="list-style-type: none"> ■ 3-pin removable screw terminal (power) ■ 4-pin removable screw terminal (relay) 	
Physical	Enclosure	Metal shell with solid mounting kits	
	Protection Class	IP30	
	Installation	1U 19" rack mount	
	Dimensions (W x H x D)	442 x 44 x 352 mm (17.4" x 1.73" x 13.85")	
LED Display	System LED	SYS, Power 1, Power 2, CFG, ALM	
	Port LED	Speed, Link, Activity	
Environment	Operating Temperature	-40°C ~ 85°C (-40°F ~ 185°F)	
	Storage Temperature	-40°C ~ 85°C (-40°F ~ 185°F)	
	Ambient Relative Humidity	10 ~ 95% (non-condensing)	
Switch Properties	MAC Address	16K-entry	
Power	Power Consumption	<ul style="list-style-type: none"> ■ 19.24W @ 110V_{AC} 	
	Power Input	<ul style="list-style-type: none"> ■ 90 ~ 264AC/88~370V_{DC} 	
Certifications	EMI	CE FCC EN55022 Class A	
	EMS	<ul style="list-style-type: none"> ■ EN 61000-4-2 ■ EN 61000-4-3 ■ EN 61000-4-4 ■ EN 61000-4-5 ■ EN 61000-4-6 ■ EN 61000-4-8 	
		Shock	IEC 60068-2-27
		Freefall	IEC 60068-2-32
		Vibration	IEC 60068-2-6

1.2 Hardware Views

1.2.1 Front View




Figure 1.1 Front View

No.	Item	Description
1	System LED panel	See "System LED Panel" on page 4 for further details.
2	Console port	RJ45 port to access the managed switch's software.
3	SFP LEDs	SFP link activity LEDs, see "System LED Panel" on page 4.
4	Reset button	Button allows for system soft reset (3 sec.) or factory default reset (5 sec.).
5	USB port	4-pin (female) port for FW backup access.
6	ETH port	10GbE SFP Port x 4.
7	LNK/ACT LED	Link activity LED.
8	SPEED LED	Speed LED.
9	ETH port	RJ45 x 24.
10	ETH port	RJ45/SFP (mini-GBIC) combo x 8.

1.2.2 Rear View




Figure 1.2 Front View

No.	Item	Description
1	Terminal block PWR2	Connect cabling for power.
2	Terminal block	Connect cabling for alarms.
3	Terminal block PWR1	Connect cabling for power.

1.2.2.1 System LED Panel




Figure 1.3 System LED Panel

No.	LED Name	LED Color	Description
1	CFG	Yellow on	TBD
		Blink yellow (1Hz)	Configuration changed, but unsaved.
		Blink yellow (3Hz)	TBD
		Blink yellow (5Hz)	TBD
		Off	Configuration saved.
2	SYS	Green on	When the EKI switch system ready.
		Blink green (1Hz)	When EKI switch system starts up.
		Blink green (3Hz)	TBD
		Blink green (5Hz)	TBD
		Off	Power on processing in uboot mode.
3	ALM	Red on	Defined major policies are detected.
		Blink red (1Hz)	Defined minor policies are detected.
		Blink red (3Hz)	TBD
		Blink red (5Hz)	TBD
		Off	Power off or system alarm is cleared or masked.
4	PWR1	Green on	Power is being supplied to power input PWR1.
		Off	Power is not being supplied to power input PWR2.
5	PWR2	Green on	Power is being supplied to power input PWR1.
		Off	Power is not being supplied to power input PWR1.
6	DATA	Green on	Link 1G
		Blink green	ACT 1G
		Amber on	Link 10/100MB
		Blink amber	ACT 10/100MB
		Off	Link down

1.2.3 Dimensions

Unit: mm




Figure 1.4 Dimensions

Chapter 2

Switch Installation

2.1 Warnings

Warning: Before working on equipment that is connected to power lines, remove any jewelry (including rings, necklaces, and watches). Metal objects can heat up when connected to power and ground, which can cause serious burns or weld the metal object to the terminals.

Caution! *Exposure to chemicals can degrade the sealing properties of materials used in the sealed relay device.*



Caution! *It is not recommended to work on the system or connect or disconnect cables during periods of lightning activity.*



Caution! *Before performing any of the following procedures, disconnect the power source from the DC circuit.*



Caution! *Read the installation instructions before connecting the system to its power source.*



Caution! *The device must be grounded. Never defeat the ground conductor or operate the equipment in the absence of a suitably installed ground conductor.*



Caution! *This unit may have more than one power supply connection. All connections must be removed to de-energize the unit.*



Caution! *The installation, replacement, or service of the device must be Only be performed by trained and qualified personnel.*



Caution! *Ultimate disposal of this product should be handled according to local and national regulations*



Caution! *To prevent the system from overheating, do not operate it in an area that exceeds the maximum recommended ambient temperature of: 70°C (158°F).*



Caution! *If the switch is to be installed in a hazardous location, ensure that the DC power source is located away from the vicinity of the switch.*



Caution! *The installation of the equipment must comply with all national and local electrical codes.*



Caution! *Explosion Hazard-The area must be known to be nonhazardous before servicing or replacing any components.*



Warning! *Airflow around the switch must be unrestricted. To prevent the switch from overheating, there must be the following minimum clearances:*



- Top and bottom: 2.0 in. (50.8 mm)
- Sides: 2.0 in. (50.8 mm)
- Front: 2.0 in. (50.8 mm)

2.2 Installation Guidelines

The following guidelines are provided to optimize the device performance. Review the guidelines before installing the device.

- Make sure cabling is away from sources of electrical noise. Radios, power lines, and fluorescent lighting fixtures can interfere with the device performance.
- Make sure the cabling is positioned away from equipment that can damage the cables.
- Operating environment is within the ranges listed range, see “Specifications” on page 2.
- Relative humidity around the switch does not exceed 95 percent (non condensing).
- Altitude at the installation site is not higher than 10,000 feet.
- In 10/100 and 10/100/1000 fixed port devices, the cable length from the switch to connected devices can not exceed 100 meters (328 feet).
- Make sure airflow around the switch and respective vents are unrestricted. Without proper airflow the switch can overheat. To prevent performance degradation and damage to the switch, make sure there is clearance at the top and bottom and around the exhaust vents.

2.3 Environment and Enclosure Guidelines

Review these environmental and enclosure guidelines before installation:

This equipment is intended for use in a Pollution Degree 2 industrial environment, in overvoltage Category II applications (as defined in IEC publication 60664-1), at altitudes up to 9842 ft (3 km) without derating.

This equipment is considered Group 1, Class A industrial equipment, according to IEC/CISPR Publication 11. Without appropriate precautions, there may be potential difficulties ensuring electromagnetic compatibility in other environments due to conducted as well as radiated disturbance.

This equipment is supplied as open-type equipment. It must be mounted within an enclosure that is suitably designed for those specific environmental conditions that will be present and appropriately designed to prevent personal injury resulting from accessibility to live parts. The enclosure must have suitable flame-retardant properties to prevent or minimize the spread of flame, complying with a flame-spread rating of 5VA, V2, V1, V0 (or equivalent) if nonmetallic. The interior of the enclosure must be accessible only by the use of a tool. Subsequent sections of this publication might contain additional information regarding specific enclosure-type ratings that are required to comply with certain product safety certifications.

2.3.1 Connecting Hardware

These instructions explain how to find a proper location for your Modbus Gateways, and how to connect to the network, hook up the power cable, and connect to the EKI-9228G Series.

2.4 Verifying Switch Operation

Before installing the device in a rack or on a wall, power on the switch to verify that the switch passes the power-on self-test (POST). To connect the cabling to the power source see “Power Supply Installation” on page 15.

At startup (POST), the System LED blinks green, while the remaining LEDs are a solid green. Once the switch passes POST self-test, the System LED turns green. The other LEDs turn off and return to their operating status. If the switch fails POST, the System LED switches to an amber state.

After a successful self-test, power down the switch and disconnect the power cabling. The switch is now ready for installation at its final location.

2.5 Installing the Switch

2.5.1 Rack-Mounting

1. Align the rack mount brackets with the holes on the switch.
2. Secure the rack mount brackets with the provided screws.




Figure 2.1 Installing the Rack Mount Brackets

3. Align the switch with the posts on the rack cabinet.
4. Secure the switch.




Figure 2.2 Installing the Switch

2.6 Installing and Removing SFP Modules

Up to two fiber optic ports are available (depends on model) for use in the switch. Refer to the technical specifications for details.

The Gigabit Ethernet ports on the switch are 100/1000Base SFP Fiber ports, which require using the 100M or 1G mini-GBIC fiber transceivers to work properly. Avantech provides completed transceiver models for different distance requirements.

The concept behind the LC port and cable is quite straightforward. Suppose that you are connecting devices I and II; contrary to electrical signals, optical signals do not require a circuit in order to transmit data. Consequently, one of the optical lines is used to transmit data from device I to device II, and the other optical line is used to transmit data from device II to device I, for full-duplex transmission.

Remember to connect the Tx (transmit) port of device I to the Rx (receive) port of device II, and the Rx (receive) port of device I to the Tx (transmit) port of device II. If you make your own cable, we suggest labeling the two sides of the same line with the same letter (A-to-A and B-to-B, or A1-to-A2 and B1-to-B2).

2.6.1 Installing SFP Modules

To connect the fiber transceiver and LC cable, use the following guidelines:

1. Remove the dust plug from the fiber optic slot chosen for the SFP transceiver.




Figure 2.3 Removing the Dust Plug from an SFP Slot

Note! Do not remove the dust plug from the SFP slot if you are not installing the transceiver at this time. The dust plug protects hardware from dust contamination.



2. Position the SFP transceiver with the handle on top, see the following figure.
3. Locate the triangular marking in the slot and align it with the bottom of the transceiver.
4. Insert the SFP transceiver into the slot until it clicks into place.

5. Make sure the module is seated correctly before sliding the module into the slot. A click sounds when it is locked in place.




Figure 2.4 Installing an SFP Transceiver

Note! *If you are attaching fiber optic cables to the transceiver, continue with the following step. Otherwise, repeat the previous steps to install the remaining SFP transceivers in the device.*



6. Remove the protective plug from the SFP transceiver.

Note! *Do not remove the dust plug from the transceiver if you are not installing the fiber optic cable at this time. The dust plug protects hardware from dust contamination.*



7. Insert the fiber cable into the transceiver. The connector snaps into place and locks.




Figure 2.5 Attaching a Fiber Optic Cable to a Transceiver

8. Repeat the previous procedures to install any additional SFP transceivers in the switch.
The fiber port is now setup.

2.6.2 Removing SFP Modules

To disconnect an LC connector, use the following guidelines:

1. Press down and hold the locking clips on the upper side of the optic cable.
2. Pull the optic cable out to release it from the transceiver.




Figure 2.6 Removing a Fiber Optic Cable to a Transceiver

3. Hold the handle on the transceiver and pull the transceiver out of the slot.




Figure 2.7 Removing an SFP Transceiver

Note! Replace the dust plug on the slot if you are not installing a transceiver. The dust plug protects hardware from dust contamination.



2.7 Connecting the Switch to Ethernet Ports

2.7.1 RJ45 Ethernet Cable Wiring

For RJ45 connectors, data-quality, twisted pair cabling (rated CAT5 or better) is recommended. The connector bodies on the RJ45 Ethernet ports are metallic and connected to the GND terminal. For best performance, use shielded cabling. Shielded cabling may be used to provide further protection.

Straight-thru Cable Wiring		Cross-over Cable Wiring	
Pin 1	Pin 1	Pin 1	Pin 3
Pin 2	Pin 2	Pin 2	Pin 6
Pin 3	Pin 3	Pin 3	Pin 1
Pin 6	Pin 6	Pin 6	Pin 2




Figure 2.8 Ethernet Plug & Connector Pin Position

Maximum cable length: 100 meters (328 ft.) for 10/100/1000BaseT.

2.8 Connecting the Switch to Console Port

The industrial switch supports a secondary means of management. By connecting the RJ45 to RS232 serial cable between a COM port on your PC (9-pin D-sub female) and the switch's RJ45 (RJ45) port, a wired connection for management can be established.




Figure 2.9 Serial Console Cable




Figure 2.10 DB 9 Pin Position

DB9 Connector	RJ45 Connector
NC	1 Orange/White
NC	2 Orange
2	3 Green/White
NC	4 Blue
5	5 Blue/White
3	6 Green
NC	7 Brown/White
NC	8 Brown




Figure 2.11 Pin Assignment

2.9 Power Supply Installation

2.9.1 Overview

Warning! Power down and disconnect the power cord before servicing or wiring the switch.



Caution! Do not disconnect modules or cabling unless the power is first switched off.



The device only supports the voltage outlined in the type plate. Do not use any other power components except those specifically designated for the switch device.

Caution! Disconnect the power cord before installation or cable wiring.



The switches can be powered using the same DC source used to power other devices. A DC voltage range of 12 to 48 VDC (Non PoE) or 48 VDC (PoE) must be applied between the V1+ terminal and the V1- terminal (PW1), see the following illustrations. A Class 2 power supply is required to maintain a UL60950 panel listing. The chassis ground screw terminal should be tied to the panel or chassis ground. A

redundant power configuration is supported through a secondary power supply unit to reduce network down time as a result of power loss.

Dual power inputs are supported and allow you to connect a backup power source.

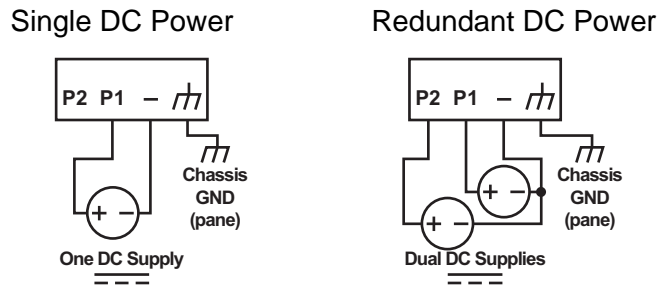


Figure 2.12 Power Wiring for EKI-9228G Series

2.9.2 Considerations

Take into consideration the following guidelines before wiring the device:

- The Terminal Block (CN1) is suitable for 12-24 AWG (3.31 - 0.205 mm²). Torque value 7 lb-in.
- The cross sectional area of the earthing conductors shall be at least 3.31 mm².
- Calculate the maximum possible current for each power and common wire. Make sure the power draw is within limits of local electrical code regulations.
- For best practices, route wiring for power and devices on separate paths.
- Do not bundle together wiring with similar electrical characteristics.
- Make sure to separate input and output wiring.
- Label all wiring and cabling to the various devices for more effective management and servicing.

Note! *Routing communications and power wiring through the same conduit may cause signal interference. To avoid interference and signal degradation, route power and communications wires through separate conduits.*



2.9.3 Grounding the Device

Caution! *Do not disconnect modules or cabling unless the power is first switched off.*



The device only supports the voltage outlined in the type plate. Do not use any other power components except those specifically designated for the switch device.

Caution! *Before connecting the device properly ground the device. Lack of a proper grounding setup may result in a safety risk and could be hazardous.*



Caution! Do not service equipment or cables during periods of lightning activity.



Caution! Do not service any components unless qualified and authorized to do so.



Caution! Do not block air ventilation holes.



2.9.4 Wiring a Relay Contact

The following section details the wiring of the relay output. The terminal block on the EKI-9728G Series is wired and then installed onto the terminal receptor located on the EKI-9228G Series.




Figure 2.13 Terminal Receptor: Relay Contact

The terminal receptor includes a total of four pins: two for PWR1 and two for PWR2.

2.9.5 Wiring the Power Inputs

Caution! Do not disconnect modules or cabling unless the power is first switched off.



The device only supports the voltage outlined in the type plate. Do not use any other power components except those specifically designated for the switch device. The temperature rating of the Input Connection Cable must be higher than 90°C.

Warning! Power down and disconnect the power cord before servicing or wiring the switch.



There are two power inputs for normal and redundant power configurations. The power input 2 is used for wiring a redundant power configuration. See the following for terminal block connector views.




Figure 2.14 Terminal Receptor: Power Input Contacts

To wire the power inputs:

Make sure the power is not connected to the switch or the power converter before proceeding.

1. Insert a small flat-bladed screwdriver in the V1+/V1- wire-clamp screws, and loosen the screws.
2. Insert the negative/positive DC wires into the V+/V- terminals of PW1. If setting up power redundancy, connect PW2 in the same manner.
3. Tighten the wire-clamp screws to secure the DC wires in place.




Figure 2.15 Installing DC Wires in a Terminal Block

2.10 Reset Button

Reset configuration to factory default:

Press and hold Reset button for 5 seconds.

System reboot:

Press and hold Reset button for 3 seconds.

Note! Do NOT power off the Ethernet switch when loading default settings.



Chapter 3

Configuration Utility

3.1 First Time Setup

3.1.1 Overview

The Industrial Ethernet Managed Switch is a configurable device that facilitates the interconnection of Ethernet devices on an Ethernet network. This includes computers, operator interfaces, I/O, controllers, RTUs, PLCs, other switches/hubs or any device that supports the standard IEEE 802.3 protocol.

This switch has all the capabilities of a store and forward Ethernet switch plus advanced management features such as SNMP, RSTP and port mirroring. This manual details how to configure the various management parameters in this easy to use switch.

3.1.2 Introduction

To take full advantage of all the features and resources available from the switch, it must be configured for your network.

The switch implements Rapid Spanning Tree Protocol (RSTP) and Simple Network Management Protocol (SNMP) to provide most of the services offered by the switch. Rapid Spanning Tree Protocol allows managed switches to communicate with each other to ensure that there exists only one active route between each pair of network nodes and provides automatic failover to the next available redundant route. A brief explanation of how RSTP works is given in the Spanning Tree section.

The switch is capable of communicating with other SNMP capable devices on the network to exchange management information. This statistical/derived information from the network is saved in the Management Information Base (MIB) of the switch. The MIB is divided into several different information storage groups. These groups will be elaborated in detail in the Management and SNMP information section of this document. The switch implements Internet Group Management Protocol (IGMP) to optimize the flow of multicast traffic on your network.

The switch supports both port-based and tag-based Virtual LANs for flexible integration with VLAN-aware networks with support for VLAN-unaware devices.

3.1.3 Administrative Interface Access

There are several administrative interfaces to the switch:

1. A graphical web interface accessible via the switch's built-in web server. Both HTTP and secure HTTPS with SSL are supported.

Note! *This is the recommended method for managing the switch.*



2. A terminal interface via the RS232/USB port or over the network using telnet or Secure Shell (SSH).
3. An SNMP interface can be used to read/write many settings.
4. Command Line Interface (CLI) can be used to read/write most settings. Initial setup must be done using an Ethernet connection (recommended) or the serial port.

3.1.4 Using the Graphical (Web) Interface

The graphical interface is provided via a web server in the switch and can be accessed via a web browser such as Opera, Mozilla, or Internet Explorer.

Note! *JavaScript must be supported and enabled in your browser for the graphical interface to work correctly.*



HTTP and HTTPS (secure HTTP) are supported for access to the web server. By default, both protocols are enabled. Either or both may be disabled to secure the switch. (See the Remote Access Security topic in this section.)

To access the graphical interface, enter a URL like HTTP://192.168.1.1 in your browser's address bar. Replace "http" with "https" to use secure http and replace "192.168.1.1" with your switch's IP address if you've changed it from the factory default.

The web server in the switch uses a signed security certificate. When you access the server via https, you may see a warning dialog indicating that the certificate was signed by an unknown authority. This is expected and to avoid this message in the future you can choose to install the certificate on your computer.

Note! *This manual describes and depicts the web user interface in detail. The terminal interface is not specifically shown but is basically the same.*



3.1.5 Configuring the Switch for Network Access

To control and monitor the switch via the network, it must be configured with basic network settings, including an IP address and subnet mask. Refer to the quick start guide in Section 1 for how to initially access your switch.

To configure the switch for network access, select [Add Menu Address Here] to reach the System Settings menu. The settings in this menu control the switch's general network configuration.

- DHCP Enabled/Disabled: The switch can automatically obtain an IP address from a server using the Dynamic Host Configuration Protocol (DHCP). This can speed up initial set up, as the network administrator does not have to find an open IP address.
- IP Address and subnet mask configuration: The IP address for the switch can be changed to a user-defined address along with a customized subnet mask to separate subnets.

Note! *Advanced users can set the IP address to 0.0.0.0 to disable the use of an IP address for additional security. However, any features requiring an IP address (i.e., web interface, etc.) will no longer be available.*



- Default Gateway Selection: A Gateway Address is chosen to be the address of a router that connects two different networks. This can be an IP address or a Fully Qualified Domain Name (FQDN) such as "domainname.org".
- NTP Server: The IP address or domain name of an NTP (Network Time Protocol) server from which the switch may retrieve the current time at startup.

Please note that using a domain name requires that at least one domain name server be configured.

3.1.6 Configuring the Ethernet Ports

The switch comes with default port settings that should allow you to connect to the Ethernet Ports with out any necessary configuration. Should there be a need to change the name of the ports, negotiation settings or flow control settings, you can do this in the Port Configuration menu. Access this menu by selecting Setup from the Main menu, and then selecting Main Settings.

- Port Name: Each port in the managed switch can be identified with a custom name. Specify a name for each port here.
- Admin: Ports can be enabled or disabled in the managed switch. For ports that are disabled, they are virtually non-existent (not visible in terms of switch operation or spanning tree algorithm). Choose to enable or disable a port by selecting Enabled or Disabled, respectively.
- Negotiation: All copper ports and gigabit fiber ports in the managed switch are capable of auto negotiation such that the fastest bandwidth is selected. Choose to enable auto-negotiation or use fixed settings. 100Mbps Fiber ports are Fixed speed only.
- Speed/Duplex/Flow Control: The managed switch accepts three local area network Ethernet Standards. The first standard, 10BASE-T, runs 10Mbps with twisted pair Ethernet cable between network interfaces. The second local area network standard is 100BASE-T, which runs at 100Mbps over the same twisted pair Ethernet cable. Lastly, there is 100BASE-F, which enables fast Ethernet (100Mbps) over fiber.

These options are available:

- 10h–10 Mbps, Half Duplex
- 10f –10 Mbps, Full Duplex
- 100h–100 Mbps, Half Duplex
- 100f –100 Mbps, Full Duplex
- 1000f–1000 Mbps, Full Duplex

On managed switches with gigabit combination ports, those ports with have two rows, a standard row of check boxes and a row labeled “SFP” with radio buttons. The SFP setting independently sets the speed at which a transceiver will operate if one is plugged in. Otherwise, the switch will use the fixed Ethernet port and the corresponding settings for it.

Note! *When 100f is selected for the SFP of a gigabit combination port, the corresponding fixed Ethernet jack will be disabled unless it is changed back to 1000F.*



3.2 Command Line Interface Configuration

3.2.1 Introduction to Command-Line Interface (CLI)

The command-line interface (CLI) is constructed with an eye toward automation of CLI-based configuration. The interaction is modeled on that used in many Internet protocols such as Telnet, FTP, and SMTP. After each command is entered and processed, the switch will issue a reply that consists of a numeric status code and a human-readable explanation of the status.

The general format of commands is:

section parameter [value]

where:

- section is used to group parameters.
- parameter will specify the parameter within the section. For example, the network section will have parameters for DHCP, IP address, subnet mask, and default gateway.
- value is the new value of the parameter. If value is omitted, the current value is displayed.

Please note that new values will not take effect until explicitly committed.

Sections and parameter names are case sensitive (e.g., “Network” is not the same as “network”).

Note! *Any commands in the CLI Commands section of this chapter, with the exception of the global commands, must be prefaced with the name of the section they are in. For example, to change the IP address of the switch, you would type:*



network address <newIP>

3.2.2 Accessing the CLI

To access the CLI interface, establish Ethernet or serial connectivity to the switch.

To connect by Ethernet, open a command prompt window and type:

telnet <switchip> (where <switchip> is the IP address of the switch)

At the login prompt, type “cli” for the username and “admin” for the password. The switch will respond with “Managed switch configuration CLI ready”.

3.3 Web Browser Configuration

The switch has an HTML based user interface embedded in the flash memory. The interface offers an easy to use means to manage basic and advanced switch functions. The interface allows for local or remote switch configuration anywhere on the network. The interface is designed for use with [Internet Explorer (6.0), Chrome, Firefox].

3.3.1 Preparing for Web Configuration

The interface requires the installation and connection of the switch to the existing network. A PC also connected to the network is required to connect to the switch and access the interface through a web browser. Use this networking information:

- IP address: 192.168.1.1
- Subnet mask: 255.255.255.0
- Default gateway: 192.168.1.254
- User name: admin
- Password: admin

3.3.2 System Login

Once the switch is installed and connected, power on the switch. The following information guides you through the logging in process.

1. Launch your web browser on the PC.
2. In the browser's address bar, type the switch's default IP address (192.168.1.1). The login screen displays.
3. Enter the user default name and password (admin / admin).
4. Click **OK** on the login screen to log in.

The main interface displays.

Chapter 4

Managing Switch

4.1 Log In

To access the login window, connect the device to the network, see “Connecting the Ethernet Media” on page 16. Once the switch is installed and connected, power on the switch see the following procedures to log into your switch.

When the switch is first installed, the default network configuration is set to DHCP enabled. You will need to make sure your network environment supports the switch setup before connecting it to the network.

1. Launch your web browser on a computer.
2. In the browser’s address bar type in the switch’s default IP address (192.168.1.1). The login screen displays.
3. Enter the default user name and password (admin/admin) to log into the management interface. You can change the default password after you have successfully logged in.
4. Click **Log In** to enter the management interface.



Figure 4.1 Login Screen

4.2 Recommended Practices

One of the easiest things to do to help increase the security posture of the network infrastructure is to implement a policy and standard for secure management. This practice is an easy way to maintain a healthy and secure network.

After you have performed the basic configurations on your switches, the following is a recommendation which is considered best practice policy.

4.2.1 Changing Default Password

In keeping with good management and security practices, it is recommended that you change the default password as soon as the device is functioning and setup correctly. The following details the necessary steps to change the default password.

To change the password:

1. Navigate to **System > Users > Accounts**.

User Name	Access Level	Lockout Status	Password Override	Password Expiration	Contained User Group	Operational Permission
admin	Privilege-15	False	Disable		default-usergroup-name	AAA: Read, Write, Execute, Debug OSPF: Read, Write, Execute, Debug
user	Privilege-1	False	Disable			

Figure 4.2 System > Users > Accounts

2. From the User Name menu, select the Admin (default) account and click **Edit**.
3. In the **User Name** field, enter admin for this account. It is not necessary to change the user name, however, a change in the default settings increases the security settings.
4. In the **Password** field, type in the new password. Re-type the same password in the **Confirm** field.
5. Click **Submit** to change the current account settings.

Figure 4.3 Changing a Default Password

After saving all the desired settings, perform a system save (**Save Configuration**). The changes are saved.

4.3 System

4.3.1 AAA

4.3.1.1 Authentication List

Use the Authentication List Configuration page to view and configure the authentication lists used for management access and port-based (IEEE 802.1X) access to the system. An authentication list specifies which authentication method(s) to use to validate the credentials of a user who attempts to access the device. Several authentication lists are preconfigured on the system. These are default lists, and they cannot be deleted. Additionally, the List Name and Access Type settings for the default lists cannot be changed.

To access this page, click **System > AAA > Authentication List**.

List Name	Access Type	Method Options	List Type	Access Line
defaultList	Login	Local	Default	Console
networkList	Login	Local	Default	Console, Telnet, SSH
enableList	Enable	Enable, None	Default	Console, Telnet, SSH
enableNetList	Enable	Enable, Deny	Default	Console, Telnet, SSH
httpList	HTTP	Local	Default	HTTP
httpsList	HTTPS	Local	Default	HTTPS
dot1xList	Dot1x	Local	Default	Dot1x

Figure 4.4 System > AAA > Authentication List

The following table describes the items in the previous figure.

Item	Description
List Name	The name of the authentication list. This field can be configured only when adding a new authentication list.

Item	Description
Access Type	<p>The way the user accesses the system. This field can be configured only when adding a new authentication list, and only the Login and Enable access types can be selected. The access types are as follows:</p> <ul style="list-style-type: none"> ■ Login: User EXEC-level management access to the command-line interface (CLI) by using a console connection or a telnet or SSH session. Access at this level has a limited number of CLI commands available to view or configure the system. ■ Enable: Privileged EXEC-level management access to the CLI by using a console connection or a telnet or SSH session. In Privileged EXEC mode, read-write users have access to all CLI commands. ■ HTTP: Management-level access to the web-based user interface by using HTTP. ■ Dot1x: Port-based access to the network through a switch port that is controlled by IEEE 802.1X.
Method Options	<p>The method(s) used to authenticate a user who attempts to access the management interface or network. The possible methods are as follows:</p> <ul style="list-style-type: none"> ■ IAS: Uses the local Internal Authentication Server (IAS) database for 802.1X port-based authentication. ■ Deny: Denies authentication. ■ Enable: Uses the locally configured Enable password to verify the user's credentials. ■ Line: Uses the locally configured Line password to verify the user's credentials. ■ Local: Uses the ID and password in the Local User database to verify the user's credentials. ■ Radius: Sends the user's ID and password to the configured Radius server to verify the user's credentials. ■ TACACS: Sends the user's ID and password to the configured TACACS server to verify the user's credentials. ■ None: No authentication is used.
List Type	<p>The type of list, which is one of the following:</p> <ul style="list-style-type: none"> ■ Default: The list is preconfigured on the system. This type of list cannot be deleted, and only the Method Options are configurable. ■ Configured: The list has been added by a user.
Access Line	<p>The access method(s) that use the list for authentication. The settings for this field are configured on the Authentication Selection page.</p>
Refresh	<p>Click Refresh to update the screen.</p>
Add	<p>Click Add to add a new authentication list. See the following procedure.</p>
Edit	<p>Click Edit to edit the selected entries.</p>

To add a new authentication list:

Click **System > AAA > Authentication List > Add**.




Figure 4.5 System > AAA > Authentication List > Add

The following table describes the items in the previous figure.

Item	Description
Authentication Methods	
Available Methods	The authentication methods that can be used for the authentication list. Not all authentication methods are available for all lists. To set the authentication method, select the method in the Available Methods field and click the right arrow to move it into the Selected Methods field.
Selected Methods	The authentication methods currently configured for the list. When multiple methods are in this field, the order in which the methods are listed is the order in which the methods will be used to authenticate a user. If the user fails to be authenticated using the first method, the device attempts to verify the user's credentials by using the next method in the list. No authentication methods can be added after None. To remove a method from this field, select it and click the left arrow to return it to the Available Methods area.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.3.1.2 Authentication Selection

Use the Authentication List Selection page to associate an authentication list with each CLI-based access method (Console, Telnet, and SSH). Each access method has the following two authentication lists associated with it:

- Login: The authentication list to use for User EXEC-level management access to the CLI. Access at this level has a limited number of CLI commands available to view or configure the system. The options available in this menu include the default Login authentication lists as well as any user-configured Login lists.
- Enable: The authentication list to use for Privileged EXEC-level management access to the CLI. In Privileged EXEC mode, read-write users have access to all CLI commands. The options available in this menu include the default Enable authentication lists as well as any user-configured Enable lists.

To access this page, click **System > AAA > Authentication Selection**.




Figure 4.6 System > AAA > Authentication Selection

The following table describes the items in the previous figure.

Item	Description
Terminal	
Console	The Login authentication list and the Enable authentication list to apply to users who attempt to access the CLI by using a connection to the console port.
Telnet	The Login authentication list and the Enable authentication list to apply to users who attempt to access the CLI by using a Telnet session.
SSH	The Login authentication list and the Enable authentication list to apply to users who attempt to access the CLI by using a secure shell (SSH) session.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.1.3 Authorization List

Use the Authorization List page to view and configure the authorization lists for users who access the command-line interface (CLI) and for users who access the network through IEEE 802.1X-enabled ports. Authorization lists are used to determine whether a user is permitted to perform a given activity on the system or network. Several authorization lists are preconfigured on the system. These are default lists, and they cannot be deleted. Additionally, the List Name and Authorization Type settings for the default lists cannot be changed.

To access this page, click **System > AAA > Authorization List**.

List Name	Authorization Type	Method Options	List Type	Access Line
dfitCmdAuthList	Commands	None	Default	Console,Telnet,SSH
dfitExecAuthList	Exec	None	Default	Console,Telnet,SSH
networkList	Network	None	Default	Dot1x

Figure 4.7 System > AAA > Authorization List

The following table describes the items in the previous figure.

Item	Description
List Name	The name of the authorization list. This field can be configured only when adding a new authorization list.
Authorization Type	The type of authorization list, which is one of the following: <ul style="list-style-type: none"> ■ Command: Determines which CLI commands a user is permitted to issue. When command authorization is enabled, each command a user enters must be validated before the command is executed. ■ EXEC: Determines whether a user can bypass User EXEC mode and enter Privileged EXEC mode directly after a successful Login authentication. ■ Network: Determines whether the user is permitted to access various network services. This authorization type applies to port-based access (IEEE 802.1X) rather than access to the CLI.

Item	Description
Method Options	<p>The method(s) used to authorize a user's access to the device or network services. The possible methods are as follows:</p> <ul style="list-style-type: none"> ■ TACACS+: When a user issues a CLI command, the device contacts the configured TACACS+ server to verify whether the user is allowed to issue the command. If approved, the command is executed. Otherwise, the command fails. ■ RADIUS: When a user is authenticated by the RADIUS server, the device downloads a list of permitted/denied commands from the RADIUS server. The list of authorized commands that are associated with the authenticated user is cached during the user's session. If this method is selected, the authentication method for the access type must also be RADIUS. ■ Local: Uses a list stored locally on the system to determine whether the user is authorized to access the given services. ■ None: No authorization is used. If the method is None, the authorization type is effectively disabled.
List Type	<p>The type of authorization list, which is one of the following:</p> <ul style="list-style-type: none"> ■ Default: The list is preconfigured on the system. This type of list cannot be deleted, and only the Method Options are configurable. ■ Configured: The list has been added by a user.
Access Line	The access method(s) that use the list for authorization. The settings for this field are configured on the Authorization Selection page.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new authorization list. See the following procedure.
Edit	Click Edit to edit the selected entries.

To add a new authorization list:

Click **System > AAA > Authorization List > Add**.

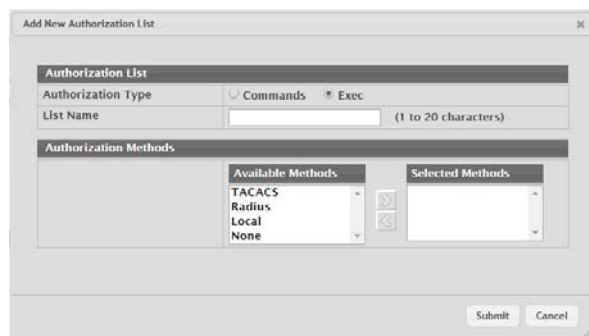


Figure 4.8 System > AAA > Authorization List > Add

The following table describes the items in the previous figure.

Item	Description
Authorization Methods	
Available Methods	The authorization methods that can be used for the authorization list. Not all methods are available for all lists. To set the authorization method, select the method in the Available Methods field and click the right arrow to move it into the Selected Methods field.

Item	Description
Selected Methods	The authorization methods currently configured for the list. When multiple methods are in this field, the order in which the methods are listed is the order in which the methods will be used to authorize a user. If the user fails to be authorized using the first method, the device attempts to authorize the user by using the next method in the list. No authorization methods can be added after None. To remove a method from this field, select it and click the left arrow to return it to the Available Methods area.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.3.1.4 Authorization Selection

Use the Authorization Selection page to associate an authorization list with each CLI-based access method (Console, Telnet, and SSH). Each access method has the following two authorization lists associated with it:

- **Exec:** The authorization list that determines whether the user is permitted to enter Privileged EXEC mode immediately after a successful Login authentication.
- **Commands:** The authorization list that determines which CLI commands the user is permitted to issue.

To access this page, click **System > AAA > Authorization Selection**.

Terminal	Access Method	Exec	Commands
Console		dfitExecAuthList	dfitCmdAuthList
Telnet		dfitExecAuthList	dfitCmdAuthList
SSH		dfitExecAuthList	dfitCmdAuthList

Submit Refresh Cancel

Figure 4.9 System > AAA > Authorization Selection

The following table describes the items in the previous figure.

Item	Description
Console	The Exec authorization list and the Commands authorization list to apply to users who access the CLI by using a connection to the console port.
Telnet	The Exec authorization list and the Commands authorization list to apply to users who access the CLI by using a Telnet session.
SSH	The Exec authorization list and the Commands authorization list to apply to users who access the CLI by using a secure shell (SSH) session.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.1.5 Accounting List

Use the Accounting List Configuration page to view and configure the accounting lists for users who access the command-line interface (CLI) to manage and monitor the device. Accounting lists are used to record user activity on the device. The device is preconfigured with accounting lists. These are default lists, and they cannot be deleted. Additionally, the List Name and Accounting Type settings for the default lists cannot be changed.

To access this page, click **System > AAA > Accounting List**.




Figure 4.10 System > AAA > Accounting List

The following table describes the items in the previous figure.

Item	Description
Accounting Type	The type of accounting list, which is one of the following: <ul style="list-style-type: none"> ■ Command: Each CLI command executed by the user, along with the time the command was executed, is recorded and sent to an external AAA server. ■ Exec: User login and logout times are recorded and sent to an external AAA server.
List Name	The name of the accounting list. This field can be configured only when adding a new accounting list.
Record Type	Indicates when to record and send information about the user activity: <ul style="list-style-type: none"> ■ StartStop: Accounting notifications are sent at the beginning and at the end of an exec session or a user-executed command. User activity does not wait for the accounting notification to be recorded at the AAA server. ■ StopOnly: Accounting notifications are sent at the end of an exec session or a user-executed command. ■ None: Accounting will not be notified.
Method Options	The method(s) used to record user activity. The possible methods are as follows: <ul style="list-style-type: none"> ■ TACACS+: Accounting notifications are sent to the configured TACACS+ server. ■ Radius: Accounting notifications are sent to the configured RADIUS server.
List Type	The type of accounting list, which is one of the following: <ul style="list-style-type: none"> ■ Default: The list is preconfigured on the system. This type of list cannot be deleted, and only the Method Options and Record Type settings are configurable. ■ Configured: The list has been added by a user.
Access Line	The access method(s) that use the list for accounting user activity. The settings for this field are configured on the Accounting Selection page.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new accounting list. See the following procedure.
Edit	Click Edit to edit the selected entries.

To add a new accounting list:

Click **System > AAA > Accounting List > Add**.

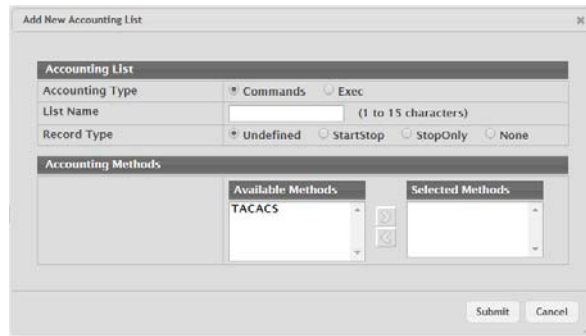


Figure 4.11 System > AAA > Accounting List > Add

The following table describes the items in the previous figure.

Item	Description
Accounting Methods	
Available Methods	The accounting methods that can be used for the accounting list. To set the accounting method, select the method in the Available Methods field and click the right arrow to move it into the Selected Methods field.
Selected Methods	The accounting methods currently configured for the list. When multiple methods are in this field, the order in which the methods are listed is the order in which the methods will be used. If the device is unable to send accounting notifications by using the first method, the device attempts to send notifications by using the second method. To remove a method from this field, select it and click the left arrow to return it to the Available Methods area.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.3.1.6 Accounting Selection

Use the Accounting List Selection page to associate an accounting list with each access method. For each access method, the following two accounting lists are associated:

- Exec: The accounting list to record user login and logout times.
- Commands: The accounting list to record which actions a user takes on the system, such as page views or configuration changes. This list also records the time when the action occurred. For Terminal access methods, this list records the CLI commands a user executes and when each command is issued.

To access this page, click **System > AAA > Accounting Selection**.




Figure 4.12 System > AAA > Accounting Selection

The following table describes the items in the previous figure.

Item	Description
Terminal	<p>The access methods in this section are CLI-based.</p> <ul style="list-style-type: none"> ■ Console: The Exec accounting list and the Commands accounting list to apply to users who access the CLI by using a connection to the console port. ■ Telnet: The Exec accounting list and the Commands accounting list to apply to users who access the CLI by using a Telnet session. ■ SSH: The Exec accounting list and the Commands accounting list to apply to users who access the CLI by using a secure shell (SSH) session.
Hypertext Transfer Protocol	<p>The access methods in this section are through a web browser.</p> <ul style="list-style-type: none"> ■ HTTP: The Exec accounting list and the Commands accounting list to apply to users who access the web-based management interface by using HTTP. ■ HTTPS: The Exec accounting list and the Commands accounting list to apply to users who access the web-based management interface by using secure HTTP (HTTPS).
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.2 Advanced Configuration

4.3.2.1 DHCP Server

Global

Use the DHCP Server Global Configuration page to configure DHCP global parameters.

To access this page, click **System > Advanced Configuration > DHCP Server > Global**.

The screenshot shows a configuration interface with the following settings:

Admin Mode	<input checked="" type="radio"/> Disable <input type="radio"/> Enable
Conflict Logging Mode	<input type="radio"/> Disable <input checked="" type="radio"/> Enable
Bootp Automatic Mode	<input checked="" type="radio"/> Disable <input type="radio"/> Enable
Ping Packet Count	<input type="text" value="2"/> (0 to 10)

Buttons: Submit, Refresh, Cancel

Figure 4.13 System > Advanced Configuration > DHCP Server > Global

The following table describes the items in the previous figure.

Item	Description
Admin Mode	Enables or disables the DHCP server administrative mode. When enabled, the device can be configured to automatically allocate TCP/IP configurations for clients.
Conflict Logging Mode	Enables or disables the logging mode for IP address conflicts. When enabled, the system stores information IP address conflicts that are detected by the DHCP server.
Bootp Automatic Mode	Enables or disables the BOOTP automatic mode. When enabled, the DHCP server supports the allocation of automatic addresses for BOOTP clients. When disabled the DHCP server supports only static addresses for BOOTP clients.

Item	Description
Ping Packet Count	The number of packets the server sends to a pool address to check for duplication as part of a ping operation. If the server receives a response to the ping, the address is considered to be in conflict and is removed from the pool.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

Excluded Addresses

Use the DHCP Server Excluded Addresses page to view and configure the IP addresses the DHCP server should not assign to clients.

To access this page, click **System > Advanced Configuration > DHCP Server > Excluded Addresses**.




Figure 4.14 System > Advanced Configuration > DHCP Server > Excluded Addresses

The following table describes the items in the previous figure.

Item	Description
From	The IP address to exclude. In a range of addresses, this value is the lowest address to exclude.
To	The highest address to exclude in a range of addresses. If the excluded address is not part of a range, this field shows the same value as the From field. When adding a single IP address to exclude, you can enter the same address specified in the From field or leave the field with the default value.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new excluded address. See the following procedure.
Remove	Click Remove to remove the selected entries.

To add a new excluded address:

Click **System > Advanced Configuration > DHCP Server > Excluded Addresses > Add**.




Figure 4.15 System > Advanced Configuration > DHCP Server > Excluded Addresses > Add

The following table describes the items in the previous figure.

Item	Description
From	The IP address to exclude. In a range of addresses, this value is the lowest address to exclude.

Item	Description
To	The highest address to exclude in a range of addresses. If the excluded address is not part of a range, this field shows the same value as the From field. When adding a single IP address to exclude, you can enter the same address specified in the From field or leave the field with the default value.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

Pool Summary

Use the DHCP Server Pool Summary page to view the currently configured DHCP server pools and to add and remove pools. A DHCP server pool is a set of network configuration information available to DHCP clients that request the information.

To access this page, click **System > Advanced Configuration > DHCP Server > Pool Summary**.




Figure 4.16 System > Advanced Configuration > DHCP Server > Pool Summary

The following table describes the items in the previous figure.

Item	Description
Name	The name that identifies the DHCP server pool.
Type	The type of binding for the pool. The options are: <ul style="list-style-type: none"> ■ Manual: The DHCP server assigns a specific IP address to the client based on the client's MAC address. This type is also known as Static. ■ Dynamic: The DHCP server can assign the client any available IP address within the pool. This type is also known as Automatic. ■ Undefined: The pool has been created by using the CLI, but the pool information has not been configured.
Network	For a Manual pool, indicates the host IP address to assign the client. For a Dynamic pool, indicates the network base address.
Lease Time	The amount of time the information the DHCP server allocates is valid.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new DHCP server pool. See the following procedure.
Remove	Click Remove to remove the selected entries.

To add a new DHCP server pool:

Click **System > Advanced Configuration > DHCP Server > Pool Summary > Add**.

Figure 4.17 System > Advanced Configuration > DHCP Server > Pool Summary > Add

The following table describes the items in the previous figure.

Item	Description
Name	The name that identifies the DHCP server pool.
Type of Binding	The type of binding for the pool. The options are: <input type="checkbox"/> Manual <input checked="" type="checkbox"/> Dynamic The binding type you select determines the fields that are available to configure.
Network Base Address	The network portion of the IP address. A DHCP client can be offered any available IP address within the defined network as long as it has not been configured as an excluded address (for dynamic pools only).
Network Mask	The subnet mask associated with the Network Base Address that separates the network bits from the host bits (for dynamic pools only).
Client Name	The system name of the client. The Client Name should not include the domain name. The function is only available for Manual pools.
Hardware Address Type	The protocol type (Ethernet or IEEE 802) used by the client's hardware platform. This value is used in response to requests from BOOTP clients. The function is only available for Manual pools.
Hardware Address	The MAC address of the client. The function is only available for Manual pools.
Client ID	The value some DHCP clients send in the Client Identifier field of DHCP messages. This value is typically identical to the Hardware Address value. In some systems, such as Microsoft DHCP clients, the client identifier is required instead of the hardware address. If the client's DHCP request includes the client identifier, the Client ID field on the DHCP server must contain the same value, and the Hardware Address Type field must be set to the appropriate value. Otherwise, the DHCP server will not respond to the client's request. The function is only available for Manual pools.
Host IP Address	The IP address to offer the client. The function is only available for Manual pools.
Host Mask	The subnet mask to offer the client. The function is only available for Manual pools.

Item	Description
Lease Expiration Mode	Indicates whether the information the server provides to the client should expire. <ul style="list-style-type: none"> ■ Enable: Allows the lease to expire. If you select this option, you can specify the amount of time the lease is valid in the Lease Duration field. ■ Disable: Sets an infinite lease time. For Dynamic bindings, an infinite lease time implies a lease period of 60 days. For a Manual binding, an infinite lease period never expires.
Lease Duration	The number of Days, Hours, and Minutes the lease is valid. This field cannot be configured if the Lease Expiration Mode is disabled.
Default Router Address	The IP address of the router to which the client should send traffic. The default router should be in the same subnet as the client. To add additional default routers, use the DHCP Server Pool Configuration page.
DNS Server Address	The IP addresses of up to two DNS servers the client should use to resolve host names into IP addresses. To add additional DNS servers, use the DHCP Server Pool Configuration page.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

Pool Configuration

Use the DHCP Server Pool Configuration page to edit pool settings or to configure additional settings for existing manual and dynamic pools. The additional settings on this page are considered advanced parameters because they are not typically used or configured. The fields that can be configured depend on the Type of Binding that is selected. The fields that do not apply to the selected binding type are disabled.






To access this page, click **System > Advanced Configuration > DHCP Server > Pool Configuration**.

Figure 4.18 System > Advanced Configuration > DHCP Server > Pool Configuration

The following table describes the items in the previous figure.

Item	Description
Pool Name	Select the pool to configure. The menu includes all pools that have been configured on the device.

Item	Description
Type of Binding	The type of binding for the pool. The options are: <ul style="list-style-type: none"> ■ Manual: The DHCP server assigns a specific IP address to the client based on the client's MAC address. This type is also known as Static. ■ Dynamic: The DHCP server can assign the client any available IP address within the pool. This type is also known as Automatic.
Network Base Address	The network portion of the IP address. A DHCP client can be offered any available IP address within the defined network as long as it has not been configured as an excluded address (for dynamic pools only).
Network Mask	The subnet mask associated with the Network Base Address that separates the network bits from the host bits (for dynamic pools only).
Client Name	The system name of the client. The Client Name should not include the domain name. This field is optional.
Hardware Address Type	The protocol type (Ethernet or IEEE 802) used by the client's hardware platform. This value is used in response to requests from BOOTP clients (for manual pools only).
Hardware Address	The MAC address of the client (for manual pools only).
Client ID	The value some DHCP clients send in the Client Identifier field of DHCP messages. This value is typically identical to the Hardware Address value. In some systems, such as Microsoft DHCP clients, the client identifier is required instead of the hardware address. If the client's DHCP request includes the client identifier, the Client ID field on the DHCP server must contain the same value, and the Hardware Address Type field must be set to the appropriate value. Otherwise, the DHCP server will not respond to the client's request (for manual pools only).
Host IP Address	The IP address to offer the client (for manual pools only).
Host Mask	For manual bindings, this field specifies the subnet mask to be statically assigned to a DHCP client. You can enter a value in Host Mask or Prefix Length to specify the subnet mask, but do not enter a value in both fields.
Lease Expiration	Indicates whether the information the server provides to the client should expire. <ul style="list-style-type: none"> ■ Enable: Allows the lease to expire. If you select this option, you can specify the amount of time the lease is valid in the Lease Duration field. ■ Disable: Sets an infinite lease time. For Dynamic bindings, an infinite lease time implies a lease period of 60 days. For a Manual binding, an infinite lease period never expires.
Lease Duration	The number of Days, Hours, and Minutes the lease is valid. This field cannot be configured if the Lease Expiration is disabled.

Item	Description
Next Server Address	<p>The IP address of the next server the client should contact in the boot process. For example, the client might be required to contact a TFTP server to download a new image file. To configure this field, click  button in the row. To reset the field to the default value, click the Reset icon in the row.</p> <p>To configure settings for one or more default routers, DNS servers, or NetBIOS servers that can be used by the client(s) in the pool, use the buttons available in the appropriate table to perform the following tasks:</p> <ul style="list-style-type: none"> ■ To add an entry to the server list, click  button and enter the IP address of the server to add. ■ To edit the address of a configured server, click  button associated with the entry to edit and update the address. ■ To delete an entry from the list, click  button associated with the entry to remove. ■ To delete all entries from the list, click  button in the heading row.
Default Router	Lists the IP address of each router to which the client(s) in the pool should send traffic. The default router should be in the same subnet as the client.
DNS Server	Lists the IP address of each DNS server the client(s) in the pool can contact to perform address resolution.
NetBIOS Server	Lists the IP address of each NetBIOS Windows Internet Naming Service (WINS) name server that is available for the selected pool.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

Pool Options

Use the DHCP Server Pool Options page to configure additional DHCP pool options, including vendor-defined options. DHCP options are collections of data with type codes that indicate how the options should be used. When a client broadcasts a request for information, the request includes the option codes that correspond to the information the client wants the DHCP server to supply.

To access this page, click **System > Advanced Configuration > DHCP Server > Pool Options**.

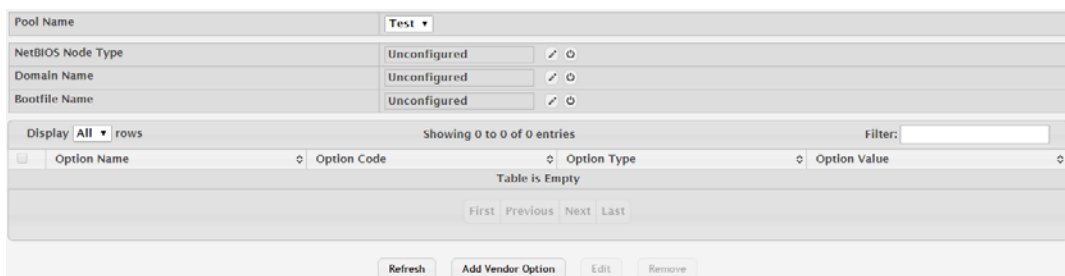


Figure 4.19 System > Advanced Configuration > DHCP Server > Pool Options
The following table describes the items in the previous figure.

Item	Description
Pool Name	Select the pool to configure. The menu includes all pools that have been configured on the device.

Item	Description
NetBIOS Node Type	The method the client should use to resolve NetBIOS names to IP addresses. To configure this field, click the Edit icon in the row. To reset the field to the default value, click the Reset icon in the row. The options are: <ul style="list-style-type: none"> ■ B-Node Broadcast: Broadcast only ■ P-Node Peer-to-Peer: NetBIOS name server only ■ M-Node Mixed: Broadcast, then NetBIOS name server ■ H-Node Hybrid: NetBIOS name server, then broadcast
Domain Name	The default domain name to configure for all clients in the selected pool.
Bootfile Name	The name of the default boot image that the client should attempt to download from a specified boot server.
Option Name	Identifies whether the entry is a fixed option or a vendor-defined option (Vendor).
Option Code	The number that uniquely identifies the option.
Option Type	The type of data to associate with the Option Code, which can be one of the following: <ul style="list-style-type: none"> ■ ASCII ■ HEX ■ IP Address
Option Value	The data associated with the Option Code. When adding or editing a vendor option, the field(s) available for configuring the value depend on the selected Option Type. If the value you configure contains characters that are not allowed by the selected Option Type, the configuration cannot be applied.
Refresh	Click Refresh to update the screen.
Add Vendor Option	Click Add Vendor Option to add a new vendor option. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

To add a new vendor option:

Click **System > Advanced Configuration > DHCP Server > Pool Options > Add Vendor Option**.

Figure 4.20 System > Advanced Configuration > DHCP Server > Pool Options > Add Vendor Option

The following table describes the items in the previous figure.

Item	Description
Option Code	The number that uniquely identifies the option.

Item	Description
Option Type	The type of data to associate with the Option Code, which can be one of the following: <ul style="list-style-type: none"> ■ ASCII ■ HEX ■ IP Address
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

To edit a vendor option:

Click **System > Advanced Configuration > DHCP Server > Pool Options > Edit**.

Figure 4.21 System > Advanced Configuration > DHCP Server > Pool Options > Edit

The following table describes the items in the previous figure.

Item	Description
Option Code	The number that uniquely identifies the option.
Option Type	The type of data to associate with the Option Code, which can be one of the following: <ul style="list-style-type: none"> ■ ASCII ■ HEX ■ IP Address
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

Bindings

Use the DHCP Server Bindings page to view information about the IP address bindings in the DHCP server database.

To access this page, click **System > Advanced Configuration > DHCP Server > Bindings**.

Figure 4.22 System > Advanced Configuration > DHCP Server > Bindings

The following table describes the items in the previous figure.

Item	Description
IP Address	The IP Address of the DHCP client.
Hardware Address	The MAC address of the DHCP client.

Item	Description
Lease Time Left	The amount of time left until the lease expires in days, hours, and minutes.
Pool Allocation Type	The type of binding used: <ul style="list-style-type: none"> ■ Dynamic: The address was allocated dynamically from a pool that includes a range of IP addresses. ■ Manual: A static IP address was assigned based on the MAC address of the client. ■ Inactive: The pool is not in use.
Refresh	Click Refresh to update the screen.
Clear Entries	Click Clear Entries to remove a selected entry.

Statistics

The DHCP Server Statistics page displays the DHCP server statistics for the device, including information about the bindings and DHCP messages. The values on this page indicate the various counts that have accumulated since they were last cleared. To access this page, click **System > Advanced Configuration > DHCP Server > Statistics**.

Automatic Bindings	0
Expired Bindings	0
Malformed Messages	0
DHCP DISCOVER packets discarded	0
Messages Received	
DHCPDISCOVER	0
DHCPREQUEST	0
DHCPDECLINE	0
DHCPRELEASE	0
DHCPINFORM	0
Messages Sent	
DHCP OFFER	0
DHCPACK	0
DHCPNAK	0

Refresh Clear Counters

Figure 4.23 System > Advanced Configuration > DHCP Server > Statistics

The following table describes the items in the previous figure.

Item	Description
Automatic Bindings	The total number of IP addresses from all address pools with automatic bindings that the DHCP server has assigned to DHCP clients.
Expired Bindings	The number of IP addresses that the DHCP server has assigned to DHCP clients that have exceeded the configured lease time.
Malformed Messages	The number of messages received from one or more DHCP clients that were improperly formatted.
DHCP DISCOVER packets discarded	The number of messages discarded from one or more DHCP Discovers.
Messages Received	
DHCPDISCOVER	The number of DHCP discovery messages the DHCP server has received. A DHCP client broadcasts this type of message to discover available DHCP servers.
DHCPREQUEST	The number of DHCP request messages the DHCP server has received. A DHCP client broadcasts this type of message in response to a DHCP offer message it received from a DHCP server.
DHCPDECLINE	The number of DHCP decline messages the DHCP server has received from clients. A client sends a decline message if the DHCP client detects that the IP address offered by the DHCP server is already in use on the network. The server then marks the address as unavailable.

Item	Description
DHCPRELEASE	The number of DHCP release messages the DHCP server has received from clients. This type of message indicates that a client no longer needs the assigned address.
DHCPINFORM	The number of DHCP inform messages the DHCP server has received from clients. A client uses this type of message to obtain DHCP options.
Messages Sent	
DHCPOFFER	The number of DHCP offer messages the DHCP server has sent to DHCP clients in response to DHCP discovery messages it has received.
DHCPACK	The number of DHCP acknowledgement messages the DHCP server has sent to DHCP clients in response to DHCP request messages it has received. The server sends this message after the client has accepted the offer from this particular server. The DHCP acknowledgement message includes information about the lease time and any other configuration information that the DHCP client has requested.
DHCPNAK	The number of negative DHCP acknowledgement messages the DHCP server has sent to DHCP clients. A server might send this type of message if the client requests an IP address that is already in use or if the server refuses to renew the lease.
Refresh	Click Refresh to update the screen.
Clear Counters	Click Clear Counters to reset all counters to zero.

Conflicts

Use the DHCP Server Conflicts Information page to view information on hosts that have address conflicts; i.e., when the same IP address is assigned to two or more devices on the network.

To access this page, click **System > Advanced Configuration > DHCP Server > Conflicts**.



Figure 4.24 System > Advanced Configuration > DHCP Server > Conflicts

The following table describes the items in the previous figure.

Item	Description
IP Address	The IP address that has been detected as a duplicate.

Item	Description
Detection Method	The method used to detect the conflict, which is one of the following: <ul style="list-style-type: none"> ■ Gratuitous ARP: The DHCP client detected the conflict by broadcasting an ARP request to the address specified in the DHCP offer message sent by the server. If the client receives a reply to the ARP request, it declines the offer and reports the conflict. ■ Ping: The server detected the conflict by sending an ICMP echo message (ping) to the IP address before offering it to the DHCP client. If the server receives a response to the ping, the address is considered to be in conflict and is removed from the pool. ■ Host Declined: The server received a DHCPDECLINE message from the host. A DHCPDECLINE message indicates that the host has discovered that the IP address is already in use on the network.
Detection Time	The time when the conflict was detected in days, hours, minutes and seconds since the system was last reset (i.e., system up time).
Refresh	Click Refresh to update the screen.
Clear Entries	Click Clear Entries to clear all of the address conflict entries.

4.3.2.2 DNS

You can use these pages to configure information about DNS servers the network uses and how the switch/ router operates as a DNS client.

Configuration

Use the DNS Global Configuration page to configure global DNS settings and to view DNS client status information.

To access this page, click **System > Advanced Configuration > DNS > Configuration**.

Figure 4.25 System > Advanced Configuration > DNS > Configuration

The following table describes the items in the previous figure.

Item	Description
Admin Mode	The administrative mode of the DNS client.
Default Domain Name	The default domain name for the DNS client to use to complete unqualified host names. Domain names are typically composed of a series of labels concatenated with dots. After a default domain name is configured, if you enter a host name and do not include the domain name information, the default domain name is automatically appended to the host name.
Retry Number	The number of times the DNS client should attempt to send DNS queries to a DNS server on the network.
Response Timeout (seconds)	The number of seconds the DNS client should wait for a response to a DNS query.

Item	Description
Domain List	The list of domain names that have been added to the DNS client's domain list. If a DNS query that includes the default domain name is not resolved, the DNS client attempts to use the domain names in this list to extend the hostname into a fully-qualified domain name. The DNS client uses the entries in the order that they appear in the list.
DNS Server	A unique IPv4 or IPv6 address used to identify a DNS server. The order in which you add servers determines the precedence of the server. The DNS server that you add first has the highest precedence and will be used before other DNS servers that you add.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

IP Mapping

Use the DNS IP Mapping page to configure DNS host names for hosts on the network and to view dynamic DNS entries. The host names are associated with IPv4 or IPv6 addresses on the network, which are statically assigned to particular hosts.

To access this page, click **System > Advanced Configuration > DNS > IP Mapping**.




Figure 4.26 System > Advanced Configuration > DNS > IP Mapping

The following table describes the items in the previous figure.

Item	Description
Entry Type	Type of DNS entry: <ul style="list-style-type: none"> ■ Static: An entry that has been manually configured on the device. ■ Dynamic: An entry that the device has learned by using a configured DNS server to resolve a hostname.
Host Name	The name that identifies the system. For Static entries, specify the Host Name after you click Add. A host name can contain up to 255 characters if it contains multiple levels in the domain hierarchy, but each level (the portion preceding a period) can contain a maximum of 63 characters. If the host name you specify is a single level (does not contain any periods), the maximum number of allowed characters is 63.
IP Address	The IPv4 or IPv6 address associated with the configured Host Name. For Static entries, specify the IP Address after you click Add. You can specify either an IPv4 or an IPv6 address.
Total Time	The number of seconds that the entry will remain in the table. The function is only available for Dynamic entries.
Elapsed Time	The number of seconds that have passed since the entry was added to the table. When the Elapsed Time reaches the Total Time, the entry times out and is removed from the table. The function is only available for Dynamic entries.
Dynamic Type	The type of address in the entry, for example IP or (less common) X.121. The function is only available for Dynamic entries.
Refresh	Click Refresh to update the screen.

Item	Description
Add	Click Add to add a new DNS entry. See the following procedure.
Remove	Click Remove to remove the selected entries.

To add a new DNS entry:

Click **System > Advanced Configuration > DNS > IP Mapping > Add**.

Figure 4.27 System > Advanced Configuration > DNS > IP Mapping > Add

The following table describes the items in the previous figure.

Item	Description
Host Name	The name that identifies the system. For Static entries, specify the Host Name after you click Add . A host name can contain up to 255 characters if it contains multiple levels in the domain hierarchy, but each level (the portion preceding a period) can contain a maximum of 63 characters. If the host name you specify is a single level (does not contain any periods), the maximum number of allowed characters is 63.
IP Address	The IPv4 or IPv6 address associated with the configured Host Name. For Static entries, specify the IP Address after you click Add . You can specify either an IPv4 or an IPv6 address.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

Source Interface Configuration

Use the DNS Source Interface Configuration page to specify the physical or logical interface to use as the DNS client source interface. When an IP address is configured on the source interface, this address is used for all DNS communications between the local DNS client and the remote DNS server. The IP address of the designated source interface is used in the IP header of DNS management protocol packets. This allows security devices, such as firewalls, to identify all source packets coming from a specific device.

To access this page, click **System > Advanced Configuration > DNS > Source Interface Configuration**.

Figure 4.28 System > Advanced Configuration > DNS > Source Interface Configuration

The following table describes the items in the previous figure.

Item	Description
Type	The type of interface to use as the source interface: <ul style="list-style-type: none"> ■ None: The primary IP address of the originating (outbound) interface is used as the source address. ■ Interface: The primary IP address of a physical port is used as the source address. ■ VLAN: The primary IP address of a VLAN routing interface is used as the source address. ■ Loopback: The primary IP address of the loopback interface is used as the source address. A loopback is always reachable, as long as any routing interface is up. ■ Tunnel: The primary IP address of a tunnel interface is used as the source address. ■ Network: The network source IP is used as the source address. ■ Service Port: The management port source IP is used as the source address.
Interface	When the selected Type is Interface, select the physical port to use as the source interface.
VLAN ID	When the selected Type is VLAN, select the VLAN to use as the source interface. The menu contains only the VLAN IDs for VLAN routing interfaces.
Loopback Interface	When the selected Type is Loopback, select the loopback interface to use as the source interface.
Tunnel ID	When the selected Type is Tunnel, select the tunnel interface to use as the source interface.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.2.3 Email Alerts

With the Email alerting feature, log messages can be sent to one or more Email addresses. You must configure information about the network Simple Mail Transport Protocol (SMTP) server for Email to be successfully sent from the switch.

The pages available from the Email Alerting folder allow you to configure information about what type of log message are sent via Email and to what address(es) the messages are delivered by Email.

Global

Use the Email Alert Global Configuration page to configure the common settings for log messages sent by the switch.

To access this page, click **System > Advanced Configuration > Email Alerts > Global**.

Figure 4.29 System > Advanced Configuration > Email Alerts > Global

The following table describes the items in the previous figure.

Item	Description
Admin Mode	Sets the administrative mode of the feature. <ul style="list-style-type: none"> ■ Enable: The device can send email alerts to the configured SMTP server. ■ Disable: The device will not send email alerts.
From Address	Specifies the email address of the sender (the switch).
Log Duration (Minutes)	Determines how frequently the non critical messages are sent to the SMTP server.
Urgent Messages Severity	Configures the severity level for log messages that are considered to be urgent. Messages in this category are sent immediately. The security level you select and all higher levels are considered to be urgent.
Non Urgent Messages Severity	Configures the severity level for log messages that are considered to be nonurgent. Messages in this category are collected and sent in a digest form at the time interval specified by the Log Duration field. The security level you select and all levels up to, but not including the lowest Urgent Messages Severity level are considered nonurgent. Messages below the security level you specify are not sent via email.
Traps Severity	The severity level for trap log messages.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

Test

Use the Email Alert Test page to verify that the Email alert settings are configured properly. After you specify the settings on this page and click **Submit**, the device will use the configured SMTP server to send an Email to the configured Email addresses. To access this page, click **System > Advanced Configuration > Email Alerts > Test**.

Figure 4.30 System > Advanced Configuration > Email Alerts > Test

The following table describes the items in the previous figure.

Item	Description
Test Message Type	Specifies the type of message to test for email alert functionality.
Test Message Body	Specifies the text contained in the body of the email alert test message.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

Server

Use the Email Alert Server Configuration page to configure information about up to three SMTP (mail) servers on the network that can handle Email alerts sent from the switch.

To access this page, click **System > Advanced Configuration > Email Alerts > Server**.




Figure 4.31 System > Advanced Configuration > Email Alerts > Server

The following table describes the items in the previous figure.

Item	Description
Address	Shows the IPv4/IPv6 address or host name of the SMTP server that handles email alerts that the device sends.
Port	Specifies the TCP port that email alerts are sent to on the SMTP server.
Security	Specifies the type of authentication to use with the mail server, which can be TLSv1 (SMTP over SSL) or None (no authentication is required).
User Name	If the Security is TLSv1, this field specifies the user name required to access the mail server.
Password	If the Security is TLSv1, this field specifies the password associated with the configured user name for mail server access. When adding or editing the server, you must retype the password to confirm that it is entered correctly.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new Email server. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

To add a new Email server:

Click **System > Advanced Configuration > Email Alerts > Server > Add**.

Figure 4.32 System > Advanced Configuration > Email Alerts > Server > Add

The following table describes the items in the previous figure.

Item	Description
Security	Specifies the type of authentication to use with the mail server, which can be TLSv1 (SMTP over SSL) or None (no authentication is required).
Port	Specifies the TCP port that Email alerts are sent to on the SMTP server.

Item	Description
User Name	If the Security is TLSv1, this field specifies the user name required to access the mail server.
Password	If the Security is TLSv1, this field specifies the password associated with the configured user name for mail server access. When adding or editing the server, you must retype the password to confirm that it is entered correctly.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

Statistics

Use the Email Alert Statistics page to view information about Email alerts sent from the switch.

To access this page, click **System > Advanced Configuration > Email Alerts > Statistics**.

Number of Emails Sent	0
Number of Emails Failed	0
Time Since Last Email Sent	0 days, 0 hours, 0 mins, 0 secs

Figure 4.33 System > Advanced Configuration > Email Alerts > Statistics

The following table describes the items in the previous figure.

Item	Description
Number of Emails Sent	The number of email alerts that were successfully sent since the counters were cleared or the system was reset.
Number of Emails Failed	The number of email alerts that failed to be sent since the counters were cleared or system was reset.
Time Since Last Email Sent	The amount of time in days, hours, minutes, and seconds that has passed since the last email alert was successfully sent.
Refresh	Click Refresh to update the screen.
Clear Counters	Click Clear Counters to reset all counters to zero.

Subject

Use the Email Alert Subject Configuration page to configure the subject line of the Email alert messages sent from the switch.

To access this page, click **System > Advanced Configuration > Email Alerts > Subject**.

Message Type	Urgent	
Email Subject	Urgent Log Messages	(1 to 255 characters)
Message Type	Email Subject	Remove
Urgent	Urgent Log Messages	<input type="checkbox"/>
Non Urgent	Non Urgent Log Messages	<input type="checkbox"/>

Figure 4.34 System > Advanced Configuration > Email Alerts > Subject

The following table describes the items in the previous figure.

Item	Description
Message Type	Select the message type with the subject to edit.
Email Subject	Specify the text to be displayed in the subject of the email alert message for the selected message type.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.

Item	Description
Delete	Click Delete to delete the selected message type.
Cancel	Click Cancel to restore default value.

Address

Use the Email Alert To Address Configuration page to configure the Email addresses to which alert messages sent.

To access this page, click **System > Advanced Configuration > Email Alerts > Address**.




Figure 4.35 System > Advanced Configuration > Email Alerts > Address

The following table describes the items in the previous figure.

Item	Description
Message Type	Specifies whether to send urgent, non urgent, or both types of email alert message to the associated address.
To Address	The valid email address of an Email alert recipient.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new Email alert to address. See the following procedure.
Remove	Click Remove to remove the selected entries.

To add a new Email alert to address:

Click **System > Advanced Configuration > Email Alerts > Address > Add**.




Figure 4.36 System > Advanced Configuration > Email Alerts > Address > Add

The following table describes the items in the previous figure.

Item	Description
To Address	The valid Email address of an Email alert recipient.
Message Type	Specifies whether to send urgent, non urgent, or both types of Email alert message to the associated address.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.3.2.4 ISDP

The Industry Standard Discovery Protocol (ISDP) is a proprietary Layer 2 network protocol which inter-operates with Cisco devices running the Cisco Discovery Protocol (CDP). ISDP is used to share information between neighboring devices. FASTPATH software participates in the CDP protocol and is able to both discover and be discovered by other CDP supporting devices.

Global

Use the ISDP Global Configuration page to configure global settings for the Industry Standard Discovery Protocol (ISDP) feature. ISDP is a proprietary Layer 2 network protocol that interoperates with the Cisco Discovery Protocol (CDP). ISDP is used to share information between neighboring devices (routers, bridges, access servers, and switches).

To access this page, click **System > Advanced Configuration > ISDP > Global**.

Figure 4.37 System > Advanced Configuration > ISDP > Global

The following table describes the items in the previous figure.

Item	Description
ISDP Mode	The administrative mode of ISDP on the device. When the mode is enabled, the device sends ISDP announcements out of each ISDP-enabled network interface that has a link partner.
ISDP V2 Mode	The administrative mode of ISDP version 2 on the device. When the mode is enabled, the device sends ISDPv2 announcements out of each ISDP-enabled network interface that has a link partner.
Message Interval (Seconds)	The number of seconds to wait between ISDP packet transmissions.
Hold Time Interval (Seconds)	The number of seconds the neighbor device should consider the information it receives in an ISDP packet to be valid.
Device ID	The identification information the device advertises to its neighbors in the ISDP packets.
Device ID Format Capability	The possible formats that the device can use for identification purposes.
Device ID Format	The current format of the device ID.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

Cache Table

Use the ISDP Cache Table page to view information about other devices the switch has discovered through the ISDP.

To access this page, click **System > Advanced Configuration > ISDP > Cache Table**.

Figure 4.38 System > Advanced Configuration > ISDP > Cache Table

The following table describes the items in the previous figure.

Item	Description
Device ID	The ID of the neighbor device as advertised in the ISDP message. The ID could be a host name, serial number, product name, MAC address, or some other type of information that identifies the neighbor device.
Interface	The local interface that is connected to the neighbor. The ISDP message was received on this interface.
IP Address	The (first) network-layer address reported in the address TLV of the most recently received ISDP message from the neighbor.
Version	The firmware version running on the neighbor device, as advertised in the ISDP message.
Hold Time (Seconds)	The number of seconds the information received in an ISDP packet is considered valid. The timer restarts each time a new ISDP packet is received from the neighbor. If the value reaches 0, the device is considered to be disconnected, and the entry ages out.
Capability	The functional capabilities advertised by the neighbor. For example, a neighbor might advertise itself as a switch, router, or host.
Platform	The hardware platform information advertised by the neighbor. The neighbor's ISDP packet might include information such as the name of the manufacturer or product model.
Port ID	The port on the neighbor device from which the ISDP packet was sent. This is the port that is directly connected to the local interface identified in the Interface field.
Protocol Version	The protocol version of the ISDP packet sent by the neighbor.
Last Time Changed	The amount of time that has passed since the entry was last modified.
Refresh	Click Refresh to update the screen.
Clear	Click Clear to remove the selected entry.

Interface

Use the ISDP Interface Configuration page to configure the ISDP settings for each interface.

To access this page, click **System > Advanced Configuration > ISDP > Interface**.



Figure 4.39 System > Advanced Configuration > ISDP > Interface

The following table describes the items in the previous figure.

Item	Description
Interface	The interface on which ISDP can be enabled or disabled. In the Edit ISDP Mode window, this field identifies the interfaces that are being configured.

Item	Description
ISDP Mode	The administrative mode of ISDP on the interface. When ISDP is enabled globally and on an interface, the interface periodically sends ISDP messages to its directly connected link partner.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.

Statistics

The ISDP Statistics page displays statistical information about the ISDP packets sent and received by the device. The transmit statistics provide information about the ISDP packets sent by all ISDP-enabled interfaces. The receive statistics provide information about the ISDP packets received from neighbor devices connected to ISDP-enabled interfaces.

To access this page, click **System > Advanced Configuration > ISDP > Statistics**.

Packets Received	0
Packets Transmitted	0
ISDPv1 Packets Received	0
ISDPv1 Packets Transmitted	0
ISDPv2 Packets Received	0
ISDPv2 Packets Transmitted	0
Bad Header	0
Checksum Error	0
Transmission Failure	0
Invalid Format Packets Received	0
Table Full	0
ISDP IP Address Table Full	0

Figure 4.40 System > Advanced Configuration > ISDP > Statistics

The following table describes the items in the previous figure.

Item	Description
Packets Received	The total number of ISDP packets received by the device.
Packets Transmitted	The total number of ISDP packets transmitted by the device.
ISDPv1 Packets Received	The total number of ISDP version 1 packets received by the device.
ISDPv1 Packets Transmitted	The total number of ISDP version 1 packets transmitted by the device.
ISDPv2 Packets Received	The total number of ISDP version 2 packets received by the device.
ISDPv2 Packets Transmitted	The total number of ISDP version 2 packets transmitted by the device.
Bad Header	The total number of ISDP packets received with bad headers.
Checksum Error	The total number of ISDP packets received with checksum errors.
Transmission Failure	The total number of ISDP packets that the device attempted to transmit but failed to do so.
Invalid Format Packets Received	The total number of ISDP packets received with an invalid ISDP packet format.
Table Full	The number of times a neighbor entry was not added to the ISDP cache table because the local database was full.
ISDP IP Address Table Full	The number of times the IP address of a neighbor could not be added to the neighbor entry because the IP address table was full.
Refresh	Click Refresh to update the screen.
Clear	Click Clear to reset all statistic to zero.

4.3.2.5 Link Dependency

The link dependency feature provides the ability to enable or disable one or more ports based on the link state of one or more different ports. With link dependency enabled on a port, the link state of that port is dependent on the link state of another port. For example, if port A is dependent on port B and the switch detects a link loss on port B, the switch automatically brings down the link on port A. When the link is restored to port B, the switch automatically restores the link to port A.

Group

Use the Link Dependency Group Status page to configure link dependency groups. Link dependency allows the link status of one interface to be dependent on the link status of another interface. Link state groups define the interface link dependency.

To access this page, click **System > Advanced Configuration > Link Dependency > Group**.




Figure 4.41 System > Advanced Configuration > Link Dependency > Group

The following table describes the items in the previous figure.

Item	Description
Group	The unique link dependency group identifier.
Downstream Interfaces	The set of interfaces dependent on other interfaces.
Upstream Interfaces	The set of interfaces that other interfaces are dependent on.
Link Action	The action performed on downstream interfaces when the upstream interfaces are down, which can be one of the following: <ul style="list-style-type: none"> ■ Up: Downstream interfaces are up when upstream interfaces are down. ■ Down: Downstream interfaces go down when upstream interfaces are down.
State	The group state, which can be one of the following: <ul style="list-style-type: none"> ■ Up: Link action is up and no upstream interfaces have their link up, or link action is down and there are upstream interfaces that have their link up. ■ Down: Link is down when the above conditions are not true.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new group. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.
Details	Click Detail to open the Group Entry Details window.

To add a new group:

Click **System > Advanced Configuration > Link Dependency > Group > Add**.

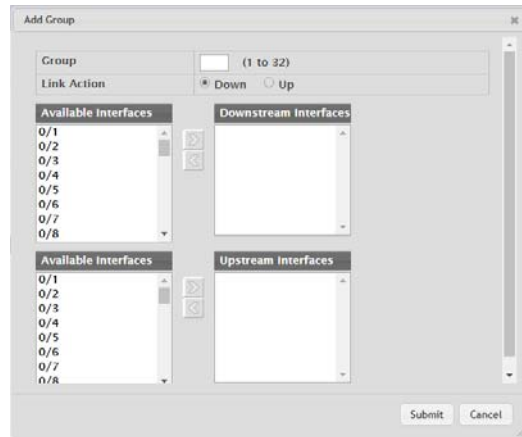


Figure 4.42 System > Advanced Configuration > Link Dependency > Group > Add

The following table describes the items in the previous figure.

Item	Description
Available Interfaces	The interfaces that can be added to the group. An interface defined as an upstream interface cannot be defined as a downstream interface in the same link state group or in a different group. Similarly, an interface defined as a downstream interface cannot be defined as an upstream interface. To move an interface between the Available Interfaces and Downstream Interfaces or Upstream Interfaces fields, click the interface (or CTRL + click to select multiple interfaces), and then click the appropriate arrow to move the selected interfaces to the desired field.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.3.2.6 LLPF

Configuration

Use the LLPF page to view and configure the Link Local Protocol Filtering (LLPF) settings on the device. The LLPF feature filters proprietary protocols that should not normally be relayed by a bridge. Using LLPF can help you troubleshoot network problems that might occur when a network includes proprietary protocols running on standards-based switches.

To access this page, click **System > Advanced Configuration > LLPF > Configuration**.

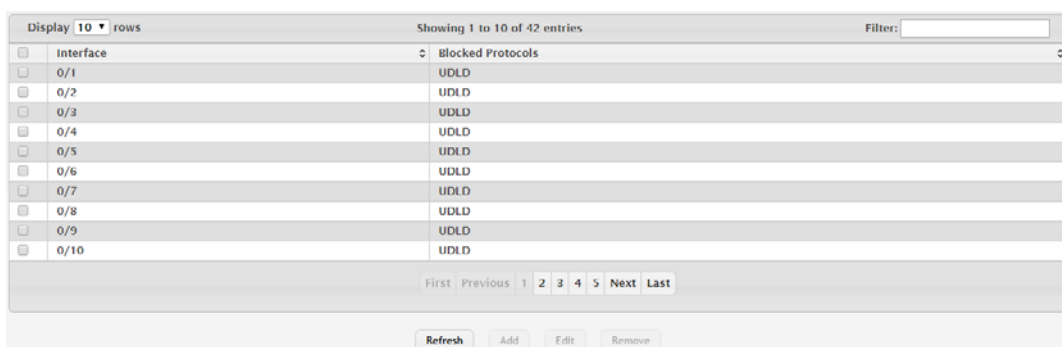


Figure 4.43 System > Advanced Configuration > LLPF > Configuration

The following table describes the items in the previous figure.

Item	Description
Interface	Identifies the physical or LAG interface.
Blocked Protocols	Lists the protocols that LLPF is configured to block on the interface. LLPF can block the following protocols: <ul style="list-style-type: none"> ■ ISDP: Industry Standard Discovery Protocol. ■ VTP: VLAN Trunking Protocol. ■ DTP: Dynamic Trunking Protocol. ■ UDLD: Uni Directional Link Detection. ■ PAGP: Port Aggregation Protocol. ■ SSTP: Shared Spanning Tree Protocol. All Protocols - All the above mentioned protocols will be dropped in addition to protocols with a Destination MAC of 01:00:0C:CC:CC:CX. When you configure the blocked protocols on the Add LLPF Interface or Edit LLPF Interface page, select the check box for each protocol to block, or clear the box to allow the protocol on the selected interface. If you select the All Protocols option, all protocols are blocked whether their associated box is checked or unchecked.
Refresh	Click Refresh to update the screen.
Add	Click Add to select the protocols for LLPF to block on an interface.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

4.3.2.7 Protection

Denial of Service

Use the Denial of Service (DoS) Configuration page to configure DoS control. FASTPATH SMB software provides support for classifying and blocking specific types of DoS attacks. You can configure your system to monitor and block these types of attacks:

To access this page, click **System > Advanced Configuration > Protection > Denial of Service**.

The screenshot displays the configuration interface for Denial of Service protection, divided into two sections: TCP Settings and ICMP Settings.

TCP Settings:

- First Fragment:
- TCP Port:
- UDP Port:
- SIP=DIP:
- SMAC=DMAC:
- TCP FIN and URG and PSH:
- TCP Flag and Sequence:
- TCP SYN:
- TCP SYN and FIN:
- TCP Fragment:
- TCP Offset:
- Min TCP Hdr Size: (0 to 255)

ICMP Settings:

- ICMP:
- Max ICMPv4 Size: (0 to 16376)
- ICMPv6:
- Max ICMPv6 Size: (0 to 16376)
- ICMP Fragment:

At the bottom of the form, there are three buttons: **Submit**, **Refresh**, and **Cancel**.

Figure 4.44 System > Advanced Configuration > Protection > Denial of Service

The following table describes the items in the previous figure.

Item	Description
TCP Settings	
First Fragment	Enable this option to allow the device to drop packets that have a TCP header smaller than the value configured in the Min TCP Hdr Size field.
TCP Port	Enable this option to allow the device to drop packets that have the TCP source port equal to the TCP destination port.
UDP Port	Enable this option to allow the device to drop packets that have the UDP source port equal to the UDP destination port.
SIP=DIP	Enable this option to allow the device to drop packets that have a source IP address equal to the destination IP address.
SMAC=DMAC	Enable this option to allow the device to drop packets that have a source MAC address equal to the destination MAC address.
TCP FIN and URG and PSH	Enable this option to allow the device to drop packets that have TCP Flags FIN, URG, and PSH set and a TCP Sequence Number equal to 0.
TCP Flag and Sequence	Enable this option to allow the device to drop packets that have TCP control flags set to 0 and the TCP sequence number set to 0.
TCP SYN	Enable this option to allow the device to drop packets that have TCP Flags SYN set.
TCP SYN and FIN	Enable this option to allow the device to drop packets that have TCP Flags SYN and FIN set.
TCP Fragment	Enable this option to allow the device to drop packets that have a TCP payload where the IP payload length minus the IP header size is less than the minimum allowed TCP header size.
TCP Offset	Enable this option to allow the device to drop packets that have a TCP header Offset set to 1.
Min TCP Hdr Size	The minimum TCP header size allowed. If First Fragment DoS prevention is enabled, the device will drop packets that have a TCP header smaller than this configured value.
ICMP Settings	
ICMP	Enable this option to allow the device to drop ICMP packets that have a type set to ECHO_REQ (ping) and a payload size greater than the ICMP payload size configured in the Max ICMPv4 Size field.
Max ICMPv4 Size	The maximum allowed ICMPv4 packet size. If ICMP DoS prevention is enabled, the device will drop ICMPv4 ping packets that have a size greater than this configured maximum ICMPv4 packet size.
ICMPv6	Enable this option to allow the device to drop ICMP packets that have a type set to ECHO_REQ (ping) and a payload size greater than the ICMP payload size configured in the Max ICMPv6 Size field.
Max ICMPv6 Size	The maximum allowed IPv6 ICMP packet size. If ICMP DoS prevention is enabled, the switch will drop IPv6 ICMP ping packets that have a size greater than this configured maximum ICMPv6 packet size.
ICMP Fragment	Enable this option to allow the device to drop fragmented ICMP packets.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.2.8 SDM

SDM

Use the SDM page to configure the Switch Database Management (SDM) template that is active after the next reboot. An SDM template is a description of the maximum resources the device can use for various features. Different SDM templates allow different combinations of scaling factors, enabling different allocations of resources depending on how the device is used. In other words, SDM templates enable you to reallocate system resources to support a different mix of features based on your network requirements.

To access this page, click **System > Advanced Configuration > SDM > SDM**.

Active SDM Template		Dual IPv4 and IPv6					
SDM Template on the Next Reload		Dual IPv4 and IPv6					
Summary							
SDM Template	ARP Entries	IPv4 Unicast Routes	IPv6 NDP Entries	IPv6 Unicast Routes	ECMP Next Hops	IPv4 Multicast Routes	IPv6 Multicast Routes
Dual IPv4 and IPv6	1536	480	512	16	1	96	32
IPv4-routing Default	2048	512	0	0	1	128	0
Data Center Plus - IPv4	2048	512	0	0	1	128	0
Data Center - IPv4 and IPv6	1536	480	512	16	1	96	32
Refresh							

Figure 4.45 System > Advanced Configuration > SDM > SDM

The following table describes the items in the previous figure.

Item	Description
Active SDM Template	The SDM template the device is currently using.
SDM Template on the Next Reload	The SDM template that will be active after the device reboots.
Summary	
ARP Entries	The maximum number of entries in the IPv4 Address Resolution Protocol (ARP) cache for routing interfaces.
IPv4 Unicast Routes	The maximum number of IPv4 unicast forwarding table entries.
IPv6 NDP Entries	The maximum number of IPv6 Neighbor Discovery Protocol (NDP) cache entries.
IPv6 Unicast Routes	The maximum number of IPv6 unicast forwarding table entries.
ECMP Next Hops	The maximum number of next hops that can be installed in the IPv4 and IPv6 unicast forwarding tables.
IPv4 Multicast Routes	The maximum number of IPv4 multicast forwarding table entries.
IPv6 Multicast Routes	The maximum number of IPv6 multicast forwarding table entries.
Refresh	Click Refresh to update the screen.

4.3.2.9 sFlow

Agent

The sFlow Agent Summary page shows information about the sFlow agent on the device. sFlow is an industry standard technology for monitoring high-speed switched and routed networks. The sFlow agent can monitor network traffic on each port and generate sFlow data to send to a centralized sFlow receiver (also known as a collector).

To access this page, click **System > Advanced Configuration > sFlow > Agent**.

Version	1.3:Broadcom Corp.:01.00.06
Agent Address	192.168.1.158
Refresh	

Figure 4.46 System > Advanced Configuration > sFlow > Agent

The following table describes the items in the previous figure.

Item	Description
Version	Identifies the version and implementation of the sFlow agent. The version string has the following structure: MIB Version; Organization; Software Version.
Agent Address	The IP address associated with the sFlow agent.
Refresh	Click Refresh to update the screen.

Receiver

Use the sFlow Receiver Configuration page to view and to edit the sFlow receiver settings. The sFlow receiver collects and analyzes information sent by the sFlow agent on the device. The sFlow agent can send packet sampling data to multiple sFlow receivers on the network.

To access this page, click **System > Advanced Configuration > sFlow > Receiver**.

Index	Owner String	Time Remaining	Maximum Datagram Size	Address	Port	Datagram Version	Monitor Session
1		0	1400	0.0.0.0	6343	5	0
2		0	1400	0.0.0.0	6343	5	0
3		0	1400	0.0.0.0	6343	5	0
4		0	1400	0.0.0.0	6343	5	0
5		0	1400	0.0.0.0	6343	5	0
6		0	1400	0.0.0.0	6343	5	0
7		0	1400	0.0.0.0	6343	5	0
8		0	1400	0.0.0.0	6343	5	0

Figure 4.47 System > Advanced Configuration > sFlow > Receiver

The following table describes the items in the previous figure.

Item	Description
Index	The receiver for which data is displayed or configured.
Owner String	The entity making use of this sFlow receiver table entry. If this field is blank, the entry is currently unclaimed.
Time Remaining	The time (in seconds) remaining before the sampler is released and stops sampling. A value of 0 essentially means the receiver is not configured. When configuring the sFlow receiver settings, you must select the Timeout Mode option before you can configure a Timeout Value.
Maximum Datagram Size	The maximum number of data bytes that can be sent in a single sample datagram. The receiver should also be set to this value to avoid fragmentation of the sFlow datagrams.
Address	The IP address of the sFlow receiver.
Port	The destination UDP port for sFlow datagrams.
Datagram Version	The version of sFlow datagrams that the sFlow agent should send to the sFlow receiver.
Monitor Session	Monitor session to enable sFlow hardware feature.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.
Clear	Click Clear to clear the selected entry.

Poller

Use the sFlow Poller Configuration page to add, remove, or edit a counter poller instance on a port (data source). Configuring a poller instance allows the sFlow agent to perform periodic counter sampling on a specified port and efficiently export counters to an sFlow receiver.

To access this page, click **System > Advanced Configuration > sFlow > Poller**.



Figure 4.48 System > Advanced Configuration > sFlow > Poller

The following table describes the items in the previous figure.

Item	Description
Poller Data Source	The sFlowDataSource for this sFlow poller. The sFlow agent supports physical ports as sFlow data sources.
Receiver Index	The sFlowReceiver for this sFlow counter poller. The specified Receiver Index must be associated with an active sFlow receiver. If a receiver expires, all pollers associated with the receiver will also expire.
Poller Interval	The maximum number of seconds between successive samples of the counters associated with this data source. A sampling interval of 0 disables counter sampling.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new poller data. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

To add a new poller data:

Click **System > Advanced Configuration > sFlow > Poller > Add**.

Figure 4.49 System > Advanced Configuration > sFlow > Poller > Add

The following table describes the items in the previous figure.

Item	Description
Poller Data Source	The sFlowDataSource for this sFlow poller. The sFlow agent supports physical ports as sFlow data sources.
Receiver Index	The sFlowReceiver for this sFlow counter poller. The specified Receiver Index must be associated with an active sFlow receiver. If a receiver expires, all pollers associated with the receiver will also expire.
Poller Interval (Seconds)	The maximum number of seconds between successive samples of the counters associated with this data source. A sampling interval of 0 disables counter sampling.

Item	Description
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

Sampler

Use the sFlow Sampler Configuration page to add, remove, or edit an sFlow sampler instance on a port (data source). Configuring a sampler instance allows the sFlow agent to perform statistical packet-based sampling of switched or routed packet flows. Packet flow sampling creates a steady, but random, stream of sFlow datagrams that are sent to the sFlow receiver.

To access this page, click **System > Advanced Configuration > sFlow > Sampler**.




Figure 4.50 System > Advanced Configuration > sFlow > Sampler

The following table describes the items in the previous figure.

Item	Description
Sampler Data Source	The sFlowDataSource for this sFlow sampler. The sFlow agent supports physical ports as sFlow data sources.
Receiver Index	The sFlowReceiver for this sFlow sampler. The specified Receiver Index must be associated with an active sFlow receiver. If a receiver expires, all samplers associated with the receiver will also expire.
Remote Agent Index	The remote agent index.
Ingress Sampling Rate	sFlow instance packet Sampling Rate for Ingress sampling.
Egress Sampling Rate	sFlow instance egress packet Sampling Rate.
Flow Based Sampling Rate	sFlow instance flow based packet Sampling Rate.
Maximum Header Size	The maximum number of bytes that should be copied from a sampled packet.
IP ACL	The ID of the IP ACL to apply to traffic from the sampler.
IP MAC	The ID of the MAC ACL to apply to traffic from the sampler.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new sampler data. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

To add a new sampler data:

Click **System > Advanced Configuration > sFlow > Sampler > Add**.

Figure 4.51 System > Advanced Configuration > sFlow > Sampler > Add

The following table describes the items in the previous figure.

Item	Description
Sampler Data Source	The sFlowDataSource for this sFlow sampler. The sFlow agent supports physical ports as sFlow data sources.
Receiver Index	The sFlowReceiver for this sFlow sampler. The specified Receiver Index must be associated with an active sFlow receiver. If a receiver expires, all samplers associated with the receiver will also expire.
Remote Agent Index	The remote agent index.
Ingress Sampling Rate	sFlow instance packet Sampling Rate for Ingress sampling.
Egress Sampling Rate	sFlow instance egress packet Sampling Rate.
Flow Based Sampling Rate	sFlow instance flow based packet Sampling Rate.
Maximum Header Size	The maximum number of bytes that should be copied from a sampled packet.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

Source Interface Configuration

Use the sFlow Source Interface Configuration page to specify the physical or logical interface to use as the sFlow client source interface. When an IP address is configured on the source interface, this address is used for all sFlow communications between the local sFlow client and the remote sFlow server. The IP address of the designated source interface is used in the IP header of sFlow management protocol packets. This allows security devices, such as firewalls, to identify all source packets coming from a specific device.

To access this page, click **System > Advanced Configuration > sFlow > Source Interface Configuration**.

Figure 4.52 System > Advanced Configuration > sFlow > Source Interface Configuration

The following table describes the items in the previous figure.

Item	Description
Type	The type of interface to use as the source interface: <ul style="list-style-type: none"> None: The primary IP address of the originating (outbound) interface is used as the source address. Interface: The primary IP address of a physical port is used as the source address. VLAN: The primary IP address of a VLAN routing interface is used as the source address. Loopback: The primary IP address of the loopback interface is used as the source address. A loopback is always reachable, as long as any routing interface is up. Tunnel: The primary IP address of a tunnel interface is used as the source address. Network: The network source IP is used as the source address. Service Port: The management port source IP is used as the source address.
Interface	When the selected Type is Interface, select the physical port to use as the source interface.
VLAN ID	When the selected Type is VLAN, select the VLAN to use as the source interface. The menu contains only the VLAN IDs for VLAN routing interfaces.
Loopback Interface	When the selected Type is Loopback, select the loopback interface to use as the source interface.
Tunnel ID	When the selected Type is Tunnel, select the tunnel interface to use as the source interface.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.2.10 SNMP

Community

Access rights are managed by defining communities on the SNMPv1, 2 Community page. When the community names are changed, access rights are also changed. SNMP Communities are defined only for SNMP v1 and SNMP v2.

Use the SNMP Community Configuration page to enable SNMP and Authentication notifications.

To access this page, click **System > Advanced Configuration > SNMP > Community**.




Figure 4.53 System > Advanced Configuration > SNMP > Community

The following table describes the items in the previous figure.

Item	Description
Community Name	Community name used in SNMPv1/v2 packets. This is configured in the client and identifies the access the user may connect with.

Item	Description
Security Name	Identifies the security entry that associates communities and Groups for a specific access type.
Group Name	Identifies the group associated with this community entry.
IP Address	Specifies the IP address that can connect with this community.
Refresh	Click Refresh to update the screen.
Add Community	Click Add Community to add a new SNMP community. See the following procedure.
Add Community Group	Click Add Community Group to add a new SNMP community group. See the following procedure.
Remove	Click Remove to remove the selected entries.

To add a new SNMP community:

Click **System > Advanced Configuration > SNMP > Community > Add Community**.

Figure 4.54 System > Advanced Configuration > SNMP > Community > Add Community

The following table describes the items in the previous figure.

Item	Description
Community Name	Community name used in SNMPv1/v2 packets. This is configured in the client and identifies the access the user may connect with.
Community Access	Specifies the access control policy for the community.
Community View	Specifies the community view for the community. If the value is empty, then no access is granted.
IP Address	Specifies the IP address that can connect with this community.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

To add a new SNMP community group:

Click **System > Advanced Configuration > SNMP > Community > Add Community Group**.

Figure 4.55 System > Advanced Configuration > SNMP > Community > Add Community Group

The following table describes the items in the previous figure.

Item	Description
Community Name	Community name used in SNMPv1/v2 packets. This is configured in the client and identifies the access the user may connect with.

Item	Description
Group Name	Identifies the Group associated with this Community entry.
IP Address	Specifies the IP address that can connect with this community.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

Trap Receiver v1/v2

Use the SNMP v1/v2 Trap Receivers page to configure settings for each SNMPv1 or SNMPv2 management host that will receive notifications about traps generated by the device. The SNMP management host is also known as the SNMP trap receiver.

To access this page, click **System > Advanced Configuration > SNMP > Trap Receiver v1/v2**.




Figure 4.56 System > Advanced Configuration > SNMP > Trap Receiver v1/v2

The following table describes the items in the previous figure.

Item	Description
Host IP Address	The IP address of the SNMP management host that will receive traps generated by the device.
Community Name	The name of the SNMP community that includes the SNMP management host and the SNMP agent on the device.
Notify Type	The type of SNMP notification to send the SNMP management host: <ul style="list-style-type: none"> ■ Inform: An SNMP message that notifies the host when a certain event has occurred on the device. The message is acknowledged by the SNMP management host. This type of notification is not available for SNMPv1. ■ Trap: An SNMP message that notifies the host when a certain event has occurred on the device. The message is not acknowledged by the SNMP management host.
SNMP Version	The version of SNMP to use, which is either SNMPv1 or SNMPv2.
Timeout Value	The number of seconds to wait for an acknowledgment from the SNMP management host before resending an inform message.
Retries	The number of times to resend an inform message that is not acknowledged by the SNMP management host.
Filter	The name of the filter for the SNMP management host. The filter is configured by using the CLI and defines which MIB objects to include or exclude from the view. This field is optional.
UDP Port	The UDP port on the SNMP management host that will receive the SNMP notifications. If no value is specified when configuring a receiver, the default UDP port value is used.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new SNMP trap receiver. See the following procedure.
Remove	Click Remove to remove the selected entries.

To add a new SNMP trap receiver:

Click **System > Advanced Configuration > SNMP > Trap Receiver v1/v2 > Add**.

Figure 4.57 System > Advanced Configuration > SNMP > Trap Receiver v1/v2 > Add

The following table describes the items in the previous figure.

Item	Description
Host IP Address	The IP address of the SNMP management host that will receive traps generated by the device.
Community Name	The name of the SNMP community that includes the SNMP management host and the SNMP agent on the device.
Notify Type	The type of SNMP notification to send the SNMP management host: <ul style="list-style-type: none"> ■ Inform: An SNMP message that notifies the host when a certain event has occurred on the device. The message is acknowledged by the SNMP management host. This type of notification is not available for SNMPv1. ■ Trap: An SNMP message that notifies the host when a certain event has occurred on the device. The message is not acknowledged by the SNMP management host.
SNMP Version	The version of SNMP to use, which is either SNMPv1 or SNMPv2.
Retries	The number of times to resend an inform message that is not acknowledged by the SNMP management host.
Timeout Value (Seconds)	The number of seconds to wait for an acknowledgment from the SNMP management host before resending an inform message.
Filter	The name of the filter for the SNMP management host. The filter is configured by using the CLI and defines which MIB objects to include or exclude from the view. This field is optional.
UDP Port	The UDP port on the SNMP management host that will receive the SNMP notifications. If no value is specified when configuring a receiver, the default UDP port value is used.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

Trap Receiver v3

Use the SNMP v3 Trap Receivers page to configure settings for each SNMPv3 management host that will receive notifications about traps generated by the device. The SNMP management host is also known as the SNMP trap receiver.

To access this page, click **System > Advanced Configuration > SNMP > Trap Receiver v3**.




Figure 4.58 System > Advanced Configuration > SNMP > Trap Receiver v3

The following table describes the items in the previous figure.

Item	Description
Host IP Address	The IP address of the SNMP management host that will receive traps generated by the device.
User Name	The name of the SNMP user that is authorized to receive the SNMP notification.
Notify Type	The type of SNMP notification to send the SNMP management host: <ul style="list-style-type: none">■ Trap: An SNMP message that notifies the host when a certain event has occurred on the device. The message is not acknowledged by the SNMP management host.■ Inform: An SNMP message that notifies the host when a certain event has occurred on the device. The message is acknowledged by the SNMP management host.
Security Level	The security level associated with the SNMP user, which is one of the following: <ul style="list-style-type: none">■ No Auth No Priv: No authentication and no data encryption (no security).■ Auth No Priv: Authentication, but no data encryption. With this security level, users send SNMP messages that use an MD5 key/password for authentication, but not a DES key/password for encryption.■ Auth Priv: Authentication and data encryption. With this security level, users send an MD5 key/password for authentication and a DES key/password for encryption.
Timeout Value	The number of seconds to wait for an acknowledgment from the SNMP receiver before resending an inform message.
Retries	The number of times to resend an inform message that is not acknowledged by the SNMP receiver.
Filter	The name of the filter for the SNMP management host. The filter is configured by using the CLI and defines which MIB objects to include or exclude from the view. This field is optional.
UDP Port	The UDP port on the SNMP management host that will receive the SNMP notifications. If no value is specified when configuring a receiver, the default UDP port value is used.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new SNMP trap receiver. See the following procedure.
Remove	Click Remove to remove the selected entries.

To add a new SNMP trap receiver:

Click **System > Advanced Configuration > SNMP > Trap Receiver v3 > Add**.

Figure 4.59 System > Advanced Configuration > SNMP > Trap Receiver v3 > Add

The following table describes the items in the previous figure.

Item	Description
Host IP Address	The IP address of the SNMP management host that will receive traps generated by the device.
User Name	The name of the SNMP user that is authorized to receive the SNMP notification.
Notify Type	The type of SNMP notification to send the SNMP management host: <ul style="list-style-type: none"> ■ Inform: An SNMP message that notifies the host when a certain event has occurred on the device. The message is acknowledged by the SNMP management host. ■ Trap: An SNMP message that notifies the host when a certain event has occurred on the device. The message is not acknowledged by the SNMP management host.
Security Level	The security level associated with the SNMP user, which is one of the following: <ul style="list-style-type: none"> ■ No Auth No Priv: No authentication and no data encryption (no security). ■ Auth No Priv: Authentication, but no data encryption. With this security level, users send SNMP messages that use an MD5 key/password for authentication, but not a DES key/password for encryption. ■ Auth Priv: Authentication and data encryption. With this security level, users send an MD5 key/password for authentication and a DES key/password for encryption.
Retries	The number of times to resend an inform message that is not acknowledged by the SNMP receiver.
Timeout Value (Seconds)	The number of seconds to wait for an acknowledgment from the SNMP receiver before resending an inform message.
Filter	The name of the filter for the SNMP management host. The filter is configured by using the CLI and defines which MIB objects to include or exclude from the view. This field is optional.
UDP Port	The UDP port on the SNMP management host that will receive the SNMP notifications. If no value is specified when configuring a receiver, the default UDP port value is used.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

Supported MIBs

The SNMP Supported MIBs page lists the MIBs that the system currently supports.

To access this page, click **System > Advanced Configuration > SNMP > Supported MIBs**.

Name	Description
RFC 1907 - SNMPv2-MIB	The MIB module for SNMPv2 entities
RFC 2819 - RMON MIB	Remote Network Monitoring Management Information Base
HC-RMON-MIB	The original version of this MIB, published as RFC3273.
HC-ALARM-MIB	Initial version of the High Capacity Alarm MIB module. This version published as RFC 3434.
HENUM-TC	A MIB module containing textual conventions for high capacity data types.
Broadcom-REF-MIB	Broadcom Reference
SNMP-COMMUNITY-MIB	This MIB module defines objects to help support coexistence between SNMPv1, SNMPv2, and SNMPv3.
SNMP-FRAMEWORK-MIB	The SNMP Management Architecture MIB
SNMP-MPD-MIB	The MIB for Message Processing and Dispatching
SNMP-NOTIFICATION-MIB	The Notification MIB Module

Figure 4.60 System > Advanced Configuration > SNMP > Supported MIBs

The following table describes the items in the previous figure.

Item	Description
Name	The RFC number, if applicable, followed by the defined name of the MIB.
Description	The RFC title, or a brief description of the MIB.
Refresh	Click Refresh to update the screen.

Access Control Group

Use the SNMP Access Control Group page to configure SNMP access control groups. These SNMP groups allow network managers to assign different levels of authorization and access rights to specific device features and their attributes. The SNMP group can be referenced by the SNMP community to provide security and context for agents receiving requests and initiating traps as well as for management systems and their tasks. An SNMP agent will not respond to a request from a management system outside of its configured group, but an agent can be a member of multiple groups at the same time to allow communication with SNMP managers from different groups. Several default SNMP groups are preconfigured on the system.

To access this page, click **System > Advanced Configuration > SNMP > Access Control Group**.

Group Name	Context Name	SNMP Version	Security Level	Read	Write	Notify
DefaultRead		SNMP V1	No Auth No Priv	Default		Default
DefaultRead		SNMP V2	No Auth No Priv	Default		Default
DefaultRead		SNMP V3	No Auth No Priv	Default		Default
DefaultRead		SNMP V3	Auth No Priv	Default		Default
DefaultRead		SNMP V3	Auth Priv	Default		Default
DefaultSuper		SNMP V1	No Auth No Priv	DefaultSuper	DefaultSuper	DefaultSuper
DefaultSuper		SNMP V2	No Auth No Priv	DefaultSuper	DefaultSuper	DefaultSuper
DefaultSuper		SNMP V3	No Auth No Priv	DefaultSuper	DefaultSuper	DefaultSuper
DefaultWrite		SNMP V1	No Auth No Priv	Default		Default
DefaultWrite		SNMP V2	No Auth No Priv	Default	Default	Default

Figure 4.61 System > Advanced Configuration > SNMP > Access Control Group

The following table describes the items in the previous figure.

Item	Description
Group Name	The name that identifies the SNMP group.

Item	Description
Context Name	The SNMP context associated with the SNMP group and its views. A user or a management application specifies the context name to get the performance information from the MIB objects associated with that context name. The Context EngineID identifies the SNMP entity that should process the request (the physical router), and the Context Name tells the agent in which context it should search for the objects requested by the user or the management application.
SNMP Version	The SNMP version associated with the group.
Security Level	The security level associated with the group, which is one of the following: <ul style="list-style-type: none"> ■ No Auth No Priv: No authentication and no data encryption (no security). This is the only Security Level available for SNMPv1 and SNMPv2 groups. ■ Auth No Priv: Authentication, but no data encryption. With this security level, users send SNMP messages that use an MD5 key/password for authentication, but not a DES key/password for encryption. ■ Auth Priv: Authentication and data encryption. With this security level, users send an MD5 key/password for authentication and a DES key/password for encryption.
Read	The level of read access rights for the group. The menu includes the available SNMP views. When adding a group, select the check box to allow the field to be configured, then select the desired view that restricts management access to viewing the contents of the agent.
Write	The level of write access rights for the group. The menu includes the available SNMP views. When adding a group, select the check box to allow the field to be configured, then select the desired view that permits management read-write access to the contents of the agent but not to the community.
Notify	The level of notify access rights for the group. The menu includes the available SNMP views. When adding a group, select the check box to allow the field to be configured, then select the desired view that permits sending SNMP traps or informs.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new access control group. See the following procedure.
Remove	Click Remove to remove the selected entries.

To add a new access control group:

Click **System > Advanced Configuration > SNMP > Access Control Group > Add**.

Figure 4.62 System > Advanced Configuration > SNMP > Access Control Group > Add

The following table describes the items in the previous figure.

Item	Description
Access Control Group	
Group Name	The name that identifies the SNMP group.
SNMP Version	The SNMP version associated with the group.
Security Level	The security level associated with the group, which is one of the following: <ul style="list-style-type: none"> ■ No Auth No Priv: No authentication and no data encryption (no security). This is the only Security Level available for SNMPv1 and SNMPv2 groups. ■ Auth No Priv: Authentication, but no data encryption. With this security level, users send SNMP messages that use an MD5 key/password for authentication, but not a DES key/password for encryption. ■ Auth Priv: Authentication and data encryption. With this security level, users send an MD5 key/password for authentication and a DES key/password for encryption.
Context Name	The SNMP context associated with the SNMP group and its views. A user or a management application specifies the context name to get the performance information from the MIB objects associated with that context name. The Context EngineID identifies the SNMP entity that should process the request (the physical router), and the Context Name tells the agent in which context it should search for the objects requested by the user or the management application.
Group Access Rights	
Read	The level of read access rights for the group. The menu includes the available SNMP views. When adding a group, select the check box to allow the field to be configured, then select the desired view that restricts management access to viewing the contents of the agent.
Write	The level of write access rights for the group. The menu includes the available SNMP views. When adding a group, select the check box to allow the field to be configured, then select the desired view that permits management read-write access to the contents of the agent but not to the community.
Notify	The level of notify access rights for the group. The menu includes the available SNMP views. When adding a group, select the check box to allow the field to be configured, then select the desired view that permits sending SNMP traps or informs.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

User Security Model

The SNMP User Security Model page provides the capability to configure the SNMP V3 user accounts.

To access this page, click **System > Advanced Configuration > SNMP > User Security Model**.




Figure 4.63 System > Advanced Configuration > SNMP > User Security Model

The following table describes the items in the previous figure.

Item	Description
User Name	Specifies the name of the SNMP user being added for the User-based Security Model (USM). Each user name must be unique within the SNMP agent user list. A user name cannot contain any leading or embedded blanks.
Group Name	A SNMP group is a group to which hosts running the SNMP service belong. A group name parameter is simply the name of that group by which SNMP communities are identified. The use of a group name provides some security and context for agents receiving requests and initiating traps and does the same for management systems and their tasks. An SNMP agent won't respond to a request from a management system outside its configured group, but an agent can be a member of multiple groups at the same time. This allows for communications with SNMP managers from different groups.
Engine ID	Each SNMPv3 agent has an engine ID that uniquely identifies the agent in the device. If given this entry will be used only for packets whose engine id is this. This field takes a hexadecimal string in the form 0102030405.
Authentication	Specifies the authentication protocol to be used on authenticated messages on behalf of the specified user. <ul style="list-style-type: none"> ■ SHA: SHA protocol will be used. ■ MD5: MD5 protocol will be used. ■ None: No authentication will be used for this user.
Privacy	Specifies the privacy protocol to be used on encrypted messages on behalf of the specified user. This parameter is only valid if the Authentication method parameter is not NONE. <ul style="list-style-type: none"> ■ DES: DES protocol will be used. ■ None: No privacy protocol will be used.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new SNMP user. See the following procedure.
Remove	Click Remove to remove the selected entries.

To add a new SNMP user:

Click **System > Advanced Configuration > SNMP > User Security Model > Add**.

Figure 4.64 System > Advanced Configuration > SNMP > User Security Model > Add

The following table describes the items in the previous figure.

Item	Description
Engine ID Type	Specifies the engine ID type to be used. <ul style="list-style-type: none"> ■ Local ■ Remote

Item	Description
Engine ID	Each SNMPv3 agent has an engine ID that uniquely identifies the agent in the device. If given this entry will be used only for packets whose engine id is this. This field takes a hexadecimal string in the form 0102030405.
User Name	Specifies the name of the SNMP user being added for the User-based Security Model (USM). Each user name must be unique within the SNMP agent user list. A user name cannot contain any leading or embedded blanks.
Group Name	A SNMP group is a group to which hosts running the SNMP service belong. A group name parameter is simply the name of that group by which SNMP communities are identified. The use of a group name provides some security and context for agents receiving requests and initiating traps and does the same for management systems and their tasks. An SNMP agent won't respond to a request from a management system outside its configured group, but an agent can be a member of multiple groups at the same time. This allows for communications with SNMP managers from different groups.
Authentication Method	Specifies the authentication protocol to be used on authenticated messages on behalf of the specified user. <ul style="list-style-type: none"> ■ SHA: SHA protocol will be used. ■ MD5: MD5 protocol will be used. ■ None: No authentication will be used for this user.
Password	Specifies the password used to generate the key to be used in authenticating messages on behalf of this user. This parameter must be specified if the Authentication method parameter is not NONE.
Privacy	Specifies the privacy protocol to be used on encrypted messages on behalf of the specified user. This parameter is only valid if the Authentication method parameter is not NONE. <ul style="list-style-type: none"> ■ DES: DES protocol will be used. ■ None: No privacy protocol will be used.
Authentication Key	Specifies the password used to generate the key to be used in encrypting messages to and from this user. This parameter must be specified if the Privacy parameter is not NONE.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

Source Interface Configuration

Use the SNMP Trap Source Interface Configuration page to specify the physical or logical interface to use as the SNMP client source interface. When an IP address is configured on the source interface, this address is used for all SNMP communications between the local SNMP client and the remote SNMP server. The IP address of the designated source interface is used in the IP header of SNMP management protocol packets. This allows security devices, such as firewalls, to identify all source packets coming from a specific device.

To access this page, click **System > Advanced Configuration > SNMP > Source Interface Configuration**.

Figure 4.65 System > Advanced Configuration > SNMP > Source Interface Configuration

The following table describes the items in the previous figure.

Item	Description
Type	The type of interface to use as the source interface: <ul style="list-style-type: none"> None: The primary IP address of the originating (outbound) interface is used as the source address. Interface: The primary IP address of a physical port is used as the source address. VLAN: The primary IP address of a VLAN routing interface is used as the source address. Loopback: The primary IP address of the loopback interface is used as the source address. A loopback is always reachable, as long as any routing interface is up. Tunnel: The primary IP address of a tunnel interface is used as the source address. Network: The network source IP is used as the source address. Service Port: The management port source IP is used as the source address.
Interface	When the selected Type is Interface, select the physical port to use as the source interface.
VLAN ID	When the selected Type is VLAN, select the VLAN to use as the source interface. The menu contains only the VLAN IDs for VLAN routing interfaces.
Loopback Interface	When the selected Type is Loopback, select the loopback interface to use as the source interface.
Tunnel ID	When the selected Type is Tunnel, select the tunnel interface to use as the source interface.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

Server Configuration

Use the SNMP Server Configuration page to view and modify the SNMP Server settings on the device. A user having sufficient privilege level may change the values shown on this page.

To access this page, click **System > Advanced Configuration > SNMP > Server Configuration**.




Figure 4.66 System > Advanced Configuration > SNMP > Server Configuration

The following table describes the items in the previous figure.

Item	Description
SNMP Server Port	The UDP port number on which the SNMP server listens for requests. Changing this value may cause existing SNMP transactions to cease communicating with the device until the client applications are reconfigured to use the new port number.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.2.11 SNTP

Global Configuration

Use the SNTP Global Configuration page to view and adjust SNTP parameters.

To access this page, click **System > Advanced Configuration > SNTP > Global Configuration**.

Client Mode	Disable
Port	None
Unicast Poll Interval (Seconds)	6 (6 to 10)
Broadcast Poll Interval (Seconds)	6 (6 to 10)
Unicast Poll Timeout (Seconds)	5 (1 to 30)
Unicast Poll Retry	1 (0 to 10)
Number of Servers Configured	None

Figure 4.67 System > Advanced Configuration > SNTP > Global Configuration

The following table describes the items in the previous figure.

Item	Description
Client Mode	Specifies the mode of operation of SNTP Client. An SNTP client may operate in one of the following modes: <ul style="list-style-type: none"> ■ Disable: SNTP is not operational. No SNTP requests are sent from the client nor are any received SNTP messages processed. ■ Unicast: SNTP operates in a point-to-point fashion. A unicast client sends a request to a designated server at its unicast address and expects a reply from which it can determine the time and, optionally the round-trip delay and local clock offset relative to the server. ■ Broadcast: SNTP operates in the same manner as multicast mode but uses a local broadcast address instead of a multicast address. The broadcast address has a single subnet scope while a multicast address has Internet wide scope.
Port	Specifies the local UDP port to listen for responses/broadcasts.
Unicast Poll Interval (Seconds)	Specifies the interval, in seconds, between unicast poll requests expressed as a power of two when configured in unicast mode.
Broadcast Poll Interval (Seconds)	Specifies the interval, in seconds, between broadcast poll requests expressed as a power of two when configured in broadcast mode. Broadcasts received prior to the expiry of this interval are discarded.
Unicast Poll Timeout (Seconds)	Specifies the timeout value, in seconds, to wait for an SNTP response when configured in unicast mode.
Unicast Poll Retry	Specifies the number of times to retry a request to an SNTP server after the first time-out before attempting to use the next configured server when configured in unicast mode.
Number of Servers Configured	Specifies the number of current valid unicast server entries configured for this client.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

Global Status

Use the SNTP Global Status page to view information about the system's SNTP client.

To access this page, click **System > Advanced Configuration > SNTP > Global Status**.

Version	4
Supported Mode	Unicast and Broadcast
Last Update Time	Jan 1 00:00:00 1970
Last Attempt Time	Jan 1 00:00:00 1970
Last Attempt Status	Other
Server IP Address	
Address Type	Unknown
Server Stratum	0
Reference Clock ID	
Server Mode	Reserved
Unicast Server Max Entries	3
Unicast Server Current Entries	0
Broadcast Count	0

Refresh

Figure 4.68 System > Advanced Configuration > SNTP > Global Status

The following table describes the items in the previous figure.

Item	Description
Version	Specifies the SNTP version the client supports.
Supported Mode	Specifies the SNTP modes the client supports. A single client can support multiple modes.
Last Update Time	Specifies the local date and time (UTC) when the SNTP client last updated the system clock.
Last Attempt Time	Specifies the local date and time (UTC) of the last SNTP request or receipt of an unsolicited message.
Last Attempt Status	<p>Specifies the status of the last SNTP request or unsolicited message for both unicast and broadcast modes. If no message has been received from a server, a status of Other is displayed. These values are appropriate for all operational modes.</p> <ul style="list-style-type: none"> ■ Other: None of the following values apply, or no message has been received. ■ Success: The SNTP operation was successful, and the system time was updated. ■ Request Timed Out: A directed SNTP request timed out without receiving a response from the SNTP server. ■ Bad Date Encoded: The time provided by the SNTP server is not valid. ■ Version Not Supported: The SNTP version supported by the server is not compatible with the version supported by the client. ■ Server Unsynchronized: The SNTP server is not synchronized with its peers. This is indicated via the leap indicator field on the SNTP message. ■ Server Kiss Of Death: The SNTP server indicated that no further queries were to be sent to this server. This is indicated by a stratum field equal to 0 in a message received from a server.
Server IP Address	Specifies the IP address or hostname of the server for the last received valid packet. If no message has been received from any server, an empty string is shown.
Address Type	Specifies the address type (IP address or DNS hostname) of the SNTP server for the last received valid packet.
Server Stratum	Specifies the claimed stratum of the server for the last received valid packet. Stratums define the accuracy of the reference clock. The higher the stratum (where zero is the highest), the more accurate the clock.

Item	Description
Reference Clock ID	Specifies the reference clock identifier of the server for the last received valid packet.
Server Mode	Specifies the mode of the server for the last received valid packet.
Unicast Server Max Entries	Specifies the maximum number of unicast server entries that can be configured on this client.
Unicast Server Current Entries	Specifies the number of current valid unicast server entries configured for this client.
Broadcast Count	Specifies the number of unsolicited broadcast SNTP messages that have been received and processed by the SNTP client since the last reboot.
Refresh	Click Refresh to update the screen.

Server Configuration

Use the SNTP Server Configuration page to view and modify information for adding and modifying Simple Network Time Protocol SNTP servers.

To access this page, click **System > Advanced Configuration > SNTP > Server Configuration**.




Figure 4.69 System > Advanced Configuration > SNTP > Server Configuration

The following table describes the items in the previous figure.

Item	Description
SNTP Server	The address or host name of an SNTP server the device can use to synchronize the system time.
Type	The configured SNTP server address type, which can be IPv4, IPv6, or DNS.
Port	The UDP port on the server to which SNTP requests are sent.
Priority	The order in which to query the servers. The SNTP client on the device continues sending SNTP requests to different servers until a successful response is received or all servers are exhausted. A server entry with a lower priority value is queried before one with a higher priority. If more than one server has the same priority, the SNTP client contacts the servers in the order that they appear in the table.
Version	Specifies the NTP version running on the server.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new SNTP server. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

To add a new SNTP server:

Click **System > Advanced Configuration > SNTP > Server Configuration > Add**.

Figure 4.70 System > Advanced Configuration > SNTP > Server Configuration > Add

The following table describes the items in the previous figure.

Item	Description
Host Name or IP Address	Specify the IPv4 address, IPv6 address, or DNS-resolvable host name of the SNTP server. Unicast SNTP requests will be sent to this address. The address you enter is displayed in the SNTP Server field on the main page. The address type is automatically detected.
Port	The UDP port on the server to which SNTP requests are sent.
Priority	The order in which to query the servers. The SNTP client on the device continues sending SNTP requests to different servers until a successful response is received or all servers are exhausted. A server entry with a lower priority value is queried before one with a higher priority. If more than one server has the same priority, the SNTP client contacts the servers in the order that they appear in the table.
Version	Specifies the NTP version running on the server.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

Server Status

The SNTP Server Status page displays status information about the SNTP servers configured on your switch.

To access this page, click **System > Advanced Configuration > SNTP > Server Status**.

Figure 4.71 System > Advanced Configuration > SNTP > Server Status

The following table describes the items in the previous figure.

Item	Description
Address	The hostname or IP address for each SNTP server that has been configured.
Last Update Time	The local date and time (UTC) included in the response from this server that was used to update the system clock.
Last Attempt Time	Specifies the local date and time (UTC) that this SNTP server was last queried.

Item	Description
Last Attempt Status	<p>Specifies the status of the last SNTP request to this server. If no packet has been received from this server, a status of Other is displayed.</p> <ul style="list-style-type: none"> ■ Other: None of the following values apply, or no message has been received. ■ Success: The SNTP operation was successful, and the system time was updated. ■ Request Timed Out: A directed SNTP request timed out without receiving a response from the SNTP server. ■ Bad Date Encoded: The time provided by the SNTP server is not valid. ■ Version Not Supported: The SNTP version supported by the server is not compatible with the version supported by the client. ■ Server Unsynchronized: The SNTP server is not synchronized with its peers. This is indicated via the leap indicator field on the SNTP message. ■ Server Kiss Of Death: The SNTP server indicated that no further queries were to be sent to this server. This is indicated by a stratum field equal to 0 in a message received from a server.
Requests	Specifies the number of SNTP requests made to this server since the system was last reset.
Failed Requests	Specifies the number of failed SNTP requests made to this server since the system was last reset.
Refresh	Click Refresh to update the screen.

Source Interface Configuration

Use the SNTP Source Interface Configuration page to specify the physical or logical interface to use as the SNTP client source interface. When an IP address is configured on the source interface, this address is used for all SNTP communications between the local SNTP client and the remote SNTP server. The IP address of the designated source interface is used in the IP header of SNTP management protocol packets. This allows security devices, such as firewalls, to identify all source packets coming from a specific device.

To access this page, click **System > Advanced Configuration > SNTP > Source Interface Configuration**.

Figure 4.72 System > Advanced Configuration > SNTP > Source Interface Configuration

The following table describes the items in the previous figure.

Item	Description
Type	The type of interface to use as the source interface: <ul style="list-style-type: none"> None: The primary IP address of the originating (outbound) interface is used as the source address. Interface: The primary IP address of a physical port is used as the source address. VLAN: The primary IP address of a VLAN routing interface is used as the source address. Loopback: The primary IP address of the loopback interface is used as the source address. A loopback is always reachable, as long as any routing interface is up. Tunnel: The primary IP address of a tunnel interface is used as the source address. Network: The network source IP is used as the source address. Service Port: The management port source IP is used as the source address.
Interface	When the selected Type is Interface, select the physical port to use as the source interface.
VLAN ID	When the selected Type is VLAN, select the VLAN to use as the source interface. The menu contains only the VLAN IDs for VLAN routing interfaces.
Loopback Interface	When the selected Type is Loopback, select the loopback interface to use as the source interface.
Tunnel ID	When the selected Type is Tunnel, select the tunnel interface to use as the source interface.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.2.12 Time Ranges

You can use these pages to configure time ranges to use in time-based access control list (ACL) rules. Time-based ACLs allow one or more rules within an ACL to be based on a periodic or absolute time. Each ACL rule within an ACL except for the implicit deny all rule can be configured to be active and operational only during a specific time period. The time range pages allow you to define specific times of the day and week in order to implement time-based ACLs. The time range is identified by a name and can then be referenced by an ACL rule defined within an ACL.

Configuration

Use the Time Range Summary page to create a named time range. Each time range can consist of one absolute time entry and/or one or more periodic time entries.

To access this page, click **System > Advanced Configuration > Time Ranges > Configuration**.




Figure 4.73 System > Advanced Configuration > Time Ranges > Configuration

The following table describes the items in the previous figure.

Item	Description
Admin Mode	Enables or disables the Time Range administrative mode. When enabled, actions with subscribed components are performed for existing time range entries.
Time Range Name	The unique ID or name that identifies this time range. A time-based ACL rule can reference the name configured in this field.
Time Range Status	Shows whether the time range is Active or Inactive. A time range is Inactive if the current day and time do not fall within any time range entries configured for the time range.
Periodic Entry Count	The number of periodic time range entries currently configured for the time range.
Absolute Entry	Shows whether an absolute time entry is currently configured for the time range.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new time range. See the following procedure.
Remove	Click Remove to remove the selected entries.

To add a new time range:

Click **System > Advanced Configuration > Time Ranges > Configuration > Add**.




Figure 4.74 System > Advanced Configuration > Time Ranges > Configuration > Add

The following table describes the items in the previous figure.

Item	Description
Time Range Name	The unique ID or name that identifies this time range. A time-based ACL rule can reference the name configured in this field.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

Entry Configuration

Use the Time Range Entry Summary page to configure entries in an existing time range configuration. Each time range configuration can have multiple Periodic entries but only one Absolute entry. A Periodic entry occurs at the same time every day or on one or more days of the week. An Absolute entry does not repeat. The start and end times for entries are based on a 24-hour clock. For example, 6:00 PM is 18:00.

Note! *The time range entries use the system time for the time periods in which they take effect. Make sure you configure the SNTP server settings so that the SNTP client on the switch can obtain the correct date and time from the server.*



To access this page, click **System > Advanced Configuration > Time Ranges > Entry Configuration**.



Figure 4.75 System > Advanced Configuration > Time Ranges > Entry Configuration

The following table describes the items in the previous figure.

Item	Description
Time Range Name	Click the drop-down menu to select a time range.
Entry Type	The type of time range entry, which is one of the following: <ul style="list-style-type: none"> ■ Absolute: Occurs once or has an undefined start or end period. The duration of an Absolute entry can be hours, days, or even years. Each time entry configuration can have only one Absolute entry. ■ Periodic: Recurring entry that takes place at fixed intervals. This type of entry occurs at the same time on one or more days of the week.
Starts	For an Absolute entry, indicates the time, day, month, and year that the entry begins. If this field is blank, the Absolute entry became active when it was configured. For a Periodic entry, indicates the time and day(s) of the week that the entry begins.
Ends	For an Absolute entry, indicates the time, day, month, and year that the entry ends. If this field is blank, the Absolute entry does not have a defined end. For a Periodic entry, indicates the time and day(s) of the week that the entry ends.
Refresh	Click Refresh to update the screen.
Add Absolute	Click Add Absolute to add a new absolute time range. See the following procedure.
Add Periodic	Click Add Periodic to add a new periodic time range. See the following procedure.
Remove	Click Remove to remove the selected entries.

To add a new absolute time range:

Click **System > Advanced Configuration > Time Ranges > Entry Configuration > Add Absolute**.

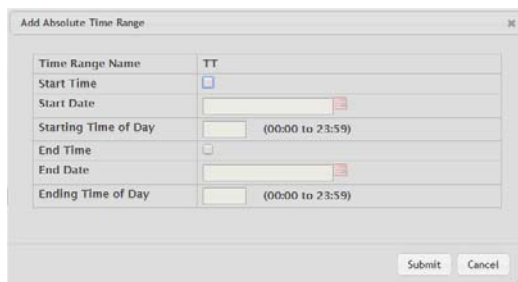


Figure 4.76 System > Advanced Configuration > Time Ranges > Entry Configuration > Add Absolute

The following table describes the items in the previous figure.

Item	Description
Time Range Name	The time range configuration that will include the Absolute time range entry.
Start Time	Select this option to configure values for the Start Date and the Starting Time of Day. If this option is not selected, the entry becomes active immediately.
Start Date	Click the calendar icon to select the day, month, and year when this entry becomes active. This field can be configured only if the Start Time option is selected.
Starting Time of Day	Specify the time of day that the entry becomes active by entering the information in the field or by using the scroll bar in the Choose Time window. Click Now to use the current time of day. Click Done to close the Choose Time window. This field can be configured only if the Start Time option is selected.
End Time	Select this option to configure values for the End Date and the Ending Time of Day. If this option is not selected, the entry does not have an end time; after the configured Start Time begins, the entry will remain active indefinitely.
End Date	Click the calendar icon to select the day, month, and year when this entry should no longer be active. This field can be configured only if the End Time option is selected.
Ending Time of Day	Specify the time of day that the entry becomes inactive by entering the information in the field or by using the scroll bar in the Choose Time window. Click Now to use the current time of day. Click Done to close the Choose Time window. This field can be configured only if the End Time option is selected.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

To add a new periodic time range:

Click **System > Advanced Configuration > Time Ranges > Entry Configuration > Add Periodic**.

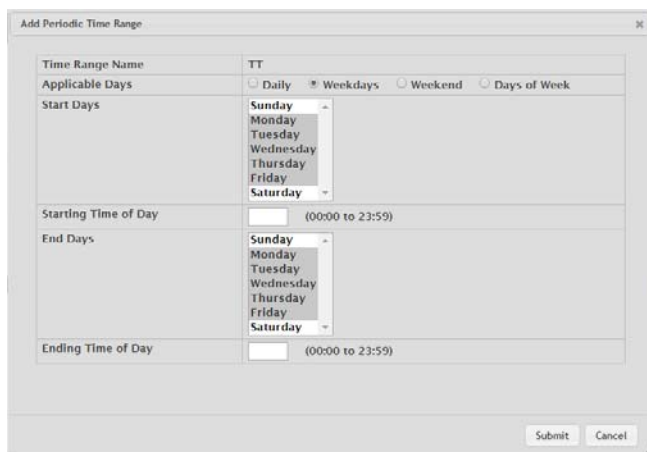


Figure 4.77 System > Advanced Configuration > Time Ranges > Entry Configuration > Add Periodic

The following table describes the items in the previous figure.

Item	Description
Time Range Name	The time range configuration that will include the Periodic time range entry.

Item	Description
Applicable Days	Select the days on which the Periodic time range entry is active: <ul style="list-style-type: none"> ■ Daily: Every day of the week ■ Weekdays: Monday through Friday ■ Weekend: Saturday and Sunday ■ Days of Week: User-defined start days
Start Days	Indicates on which days the time entry becomes active. If the selected option in the Applicable Days field is Days of Week, select one or more days on which the entry becomes active. To select multiple days, hold the CTRL key and select each desired start day.
Starting Time of Day	Specify the time of day that the entry becomes active by entering the information in the field or by using the scroll bar in the Choose Time window. Click Now to use the current time of day. Click Done to close the Choose Time window.
End Days	Indicates on which days the time entry ends. If the selected option in the Applicable Days field is Days of Week, select one or more days on which the entry ends. To select multiple days, hold the CTRL key and select each desired end day.
Ending Time of Day	Specify the time of day that the entry becomes inactive by entering the information in the field or by using the scroll bar in the Choose Time window. Click Now to use the current time of day. Click Done to close the Choose Time window.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.3.2.13 Time Zone

Summary

The Time Zone Summary page displays information about the current system time, the time zone, and the daylight saving time (also known as summer time) settings configured on the device.

To access this page, click **System > Advanced Configuration > Time Zone > Summary**.

Current Time	
Time	17:45:25
Zone	(UTC+0:00)
Date	January 01, 2016
Time Source	No time source
Time Zone	
Zone	
Offset	UTC+0:00
Summer Time	
Summer Time	No Summer Time
Zone	
Offset	
Status	

Refresh

Figure 4.78 System > Advanced Configuration > Time Zone > Summary

The following table describes the items in the previous figure.

Item	Description
Current Time	
Time	The current time on the system clock. This time is used to provide time stamps on log messages. Additionally, some CLI show commands include the time in the command output.
Zone	The acronym that represents the time zone.
Date	The current date on the system.

Item	Description
Time Source	The time source from which the time update is taken: <ul style="list-style-type: none"> ■ SNTP: The time has been acquired from an SNTP server. ■ No Time Source: The time has either been manually configured or not configured at all.
Time Zone	
Zone	The acronym that represents the time zone.
Offset	The number of hours offset from Coordinated Universal Time (UTC), which is also known as Greenwich Mean Time (GMT).
Summer Time	
Summer Time	The summer time mode on the system: <ul style="list-style-type: none"> ■ Disable: Summer time is not active, and the time does not shift based on the time of year. ■ Recurring: Summer time occurs at the same time every year. The start and end times and dates for the time shift must be manually configured. ■ EU: The system clock uses the standard recurring summer time settings used in countries in the European Union. When this field is selected, the rest of the applicable fields on the page except Offset and Zone are automatically populated and cannot be edited. ■ USA: The system clock uses the standard recurring daylight saving time settings used in the United States. When this field is selected, the rest of the applicable fields on the page except Offset and Zone are automatically populated and cannot be edited. ■ Non-Recurring: Summer time settings are in effect only between the start date and end date of the specified year. When this mode is selected, the summer time settings do not repeat on an annual basis.
Zone	The acronym that represents the time zone of the summer time.
Offset	The number of hours offset from Coordinated Universal Time (UTC), which is also known as Greenwich Mean Time (GMT).
Status	Indicates if summer time is currently active.
Refresh	Click Refresh to update the screen.

Time Zone

Use the Time Zone Configuration page to manually configure the system clock settings. The SNTP client must be disabled to allow manual configuration of the system time and date.

To access this page, click **System > Advanced Configuration > Time Zone > Time Zone**.

Figure 4.79 System > Advanced Configuration > Time Zone > Time Zone

The following table describes the items in the previous figure.

Item	Description
Time Zone	

Item	Description
Offset	The system clock's offset from UTC, which is also known as Greenwich Mean Time (GMT).
Zone	The acronym that represents the time zone. This field is not validated against an official list of time zone acronyms.
Date and Time	
Time	The current time in hours, minutes, and seconds on the system clock.
Date	The current date in month, day, and year on the system clock. To change the date, click the calendar icon to the right of the field, select the year from the menu, browse to the desired month, and click the date.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

Summer Time

Use the Summer Time Configuration page to configure settings for summer time, which is also known as daylight saving time. Used in some countries around the world, summer time is the practice of temporarily advancing clocks during the summer months. Typically clocks are adjusted forward one or more hours near the start of spring and are adjusted backward in autumn.

To access this page, click **System > Advanced Configuration > Time Zone > Summer Time**.

Figure 4.80 System > Advanced Configuration > Time Zone > Summer Time

The following table describes the items in the previous figure.

Item	Description
Summer Time	<p>The summer time mode on the system:</p> <ul style="list-style-type: none"> ■ Disable: Summer time is not active, and the time does not shift based on the time of year. ■ Recurring: Summer time occurs at the same time every year. The start and end times and dates for the time shift must be manually configured. ■ EU: The system clock uses the standard recurring summer time settings used in countries in the European Union. When this field is selected, the rest of the applicable fields on the page except Offset and Zone are automatically populated and cannot be edited. ■ USA: The system clock uses the standard recurring daylight saving time settings used in the United States. When this field is selected, the rest of the applicable fields on the page except Offset and Zone are automatically populated and cannot be edited. ■ Non-Recurring: Summer time settings are in effect only between the start date and end date of the specified year. When this mode is selected, the summer time settings do not repeat on an annual basis.
Date Range	
Start Date	The day, month, and year that summer time begins. To change the date, click the calendar icon to the right of the field, select the year from the menu, browse to the desired month, and click the date.
Starting Time of Day	The time, in hours and minutes, to start summer time on the specified day.
End Date	The day, month, and year that summer time ends. To change the date, click the calendar icon to the right of the field, select the year from the menu, browse to the desired month, and click the date.
Ending Time of Day	The time, in hours and minutes to end summer time on the specified day.
Recurring Date	
Start Week	The week of the month within which summer time begins.
Start Day	The day of the week on which summer time begins.
Start Month	The month of the year within which summer time begins.
Starting Time of Day	The time, in hours and minutes, to start summer time.
End Week	The week of the month within which summer time ends.
End Day	The day of the week on which summer time ends.
End Month	The month of the year within which summer time ends.
Ending Time of Day	The time, in hours and minutes, to end summer time.
Zone	
Offset	The number of minutes to shift the summer time from the standard time.
Zone	The acronym associated with the time zone when summer time is in effect.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.2.14 Trap Manager

The pages in the Trap Manager folder allow you to view and configure information about alarm LED, alarm relay, alarm relay2, logs, Email and SNMP traps the system generates.

Trap Log

Use the System Trap Log page to view the entries in the trap log.

To access this page, click **System > Advanced Configuration > Trap Manager > Trap Log**.

Trap Log Capacity	256
Number of Traps Since Last Reset	7
Number of Traps Since Log Last Viewed	7

Display	All	rows	Showing 1 to 7 of 7 entries	Filter:	<input type="text"/>
Log	System Up Time	Trap			
0	Jan 1 16:34:32 2016	Multiple Users: CPU			
1	Jan 1 15:53:47 2016	Multiple Users: CPU			
2	Jan 1 15:53:41 2016	Multiple Users: CPU			
3	Jan 1 15:53:36 2016	Multiple Users: CPU			
4	Jan 1 15:45:26 2016	Multiple Users: CPU			
5	Jan 1 15:40:44 2016	Cold Start: Unit: 0			
6	Jan 1 15:39:54 2016	Link Up: ge0/1			

First Previous 1 Next Last

Refresh Clear Log

Figure 4.81 System > Advanced Configuration > Trap Manager > Trap Log

The following table describes the items in the previous figure.

Item	Description
Trap Log Capacity	The maximum number of traps the log can store. If the number of traps exceeds the capacity, new entries overwrite the oldest entries.
Number of Traps Since Last Reset	The number of traps the system has generated since the trap log entries were last cleared, either by clicking Clear Log or by resetting the system.
Number of Traps Since Log Last Viewed	The number of traps the system has generated since the traps were last displayed. Displaying the traps by any available method (for example, uploading the file from the switch or viewing the logs from a terminal interface) will cause this counter to be reset to 0.
Log	The sequence number of this trap.
System Up Time	The time at which this trap occurred, expressed in days, hours, minutes and seconds since the device was last reset.
Trap	Provides information about the trap.
Refresh	Click Refresh to update the screen.
Clear Log	Click Clear Log to clear the current entries from the log file and resets the counters.

Trap Flags

Use the Trap Flags page to specify which software features should generate SNMP traps. If the trap flag is enabled for a feature and a significant event occurs, the SNMP agent on the device sends a trap message to any enabled SNMP trap receivers and writes a message to the trap log.

To access this page, click **System > Advanced Configuration > Trap Manager > Trap Flags**.

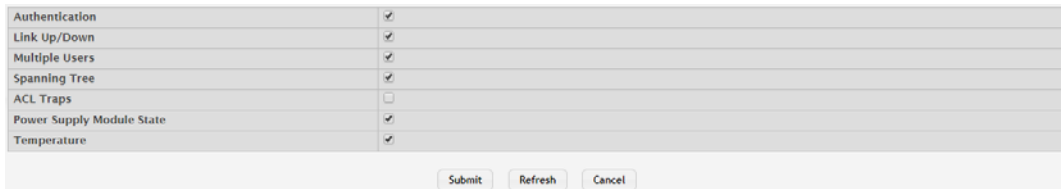


Figure 4.82 System > Advanced Configuration > Trap Manager > Trap Flags

The following table describes the items in the previous figure.

Item	Description
Authentication	Specify whether to enable SNMP notifications when events involving authentication occur, such as when a user attempts to access the device management interface and fails to provide a valid username and password.
Link Up/Down	Specify whether to enable SNMP notifications when the administrative or operational state of a physical or logical link changes.
Multiple Users	Specify whether to enable SNMP notifications when the same user ID is logged into the device more than once at the same time (either via telnet or the serial port).
Spanning Tree	Specify whether to enable SNMP notifications when various spanning tree events occur.
ACL Traps	Specify whether to enable SNMP notifications when a packet matches a configured ACL rule that includes ACL logging.
Power Supply Module State	Specify whether to enable SNMP notifications when power supply events occur.
Temperature	Specify whether to enable SNMP notifications when temperature events occur.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.2.15 CPU Traffic Filter

Global

Use the Global page to view and modify the CPU Traffic Filter settings on the device. To access this page, click **System > Advanced Configuration > CPU Traffic Filter > Global**.

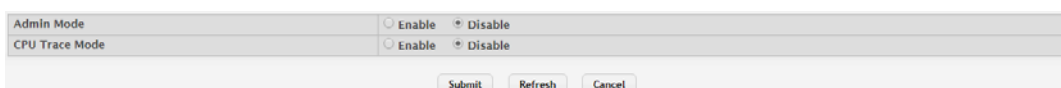


Figure 4.83 System > Advanced Configuration > CPU Traffic Filter > Global

The following table describes the items in the previous figure.

Item	Description
Admin Mode	This configures CPU-traffic mode. The packets in the Rx/Tx directions are matched when the mode is enabled. The default value is disabled.
CPU Trace Mode	This configures CPU packet tracing. The packet may be received by multiple components. If the feature is enabled and tracing configured then the packets are traced per the defined filter.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

Filter Configuration

Use the Filter Configuration page to create or remove CPU Traffic Filters and to view summary information about the filters that exist on the device.

To access this page, click **System > Advanced Configuration > CPU Traffic Filter > Filter Configuration**.

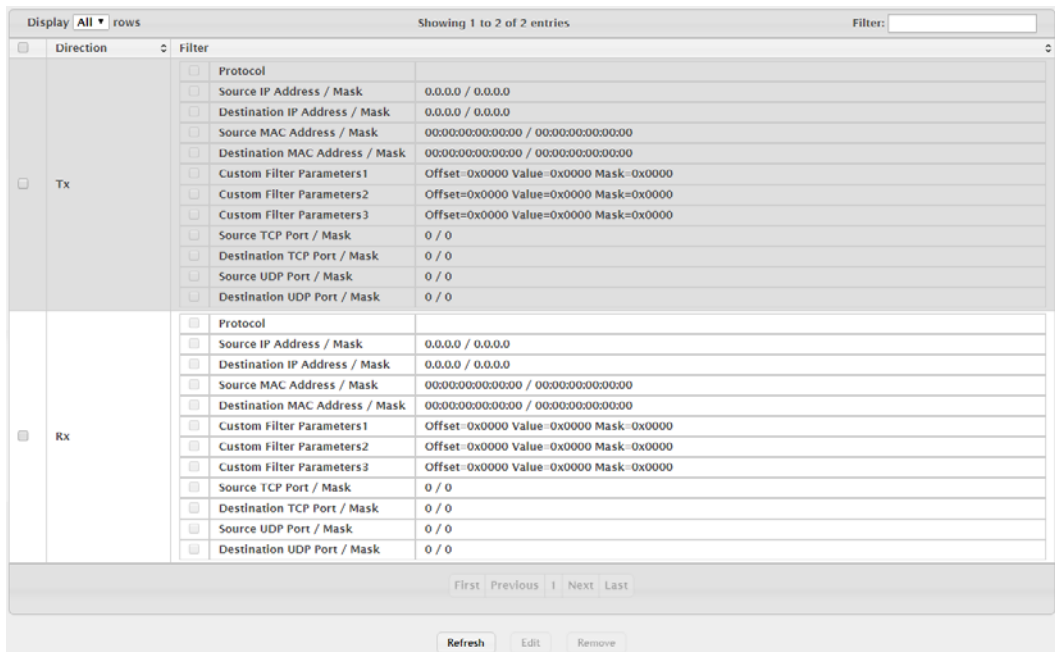


Figure 4.84 System > Advanced Configuration > CPU Traffic Filter > Filter Configuration

The following table describes the items in the previous figure.

Item	Description
Direction	Only two software filters are used (one filter each direction Tx or Rx) with condition matching as one, many or all in the below list in Tx or Rx or Both direction.
Filter	This shows specific filters for each direction.

Item	Description
Protocol	<p>This configures specific protocol filters. The statistics and/or the traces for configured filters are obtained for the packet matching configured filter. The protocol options are:</p> <ul style="list-style-type: none"> ■ STP ■ LACPDU ■ ARP ■ UDLD ■ LLDP ■ IP ■ OSPF ■ BGP ■ DHCP ■ BCAST ■ MCAST ■ UCAST ■ Source IP ■ Destination IP ■ Source MAC ■ Destination MAC ■ Custom ■ Source TCP ■ Destination TCP ■ Source UDP ■ Destination UDP
IP Address	This configures Source or Destination IP address specific filter. The statistics and/or the traces for configured filters are obtained for the packet matching configured Source / Destination IP/Mask.
MAC Address	This configures Source / Destination MAC address specific filter. The statistics and/or the traces for configured filters are obtained for the packet matching configured Source / Destination MAC address.
Custom	This configures custom filter. The statistics and/or the traces for configured filters are obtained for the packet matching configured data at specific offset. If the mask is not specified then default mask is 0xFF.
TCP Port	This configures Source / Destination TCP Port specific filter. The statistics and/or the traces for configured filters are obtained for the packet matching configured Source / Destination TCP Port.
UDP Port	This configures Source / Destination UDP Port specific filter. The statistics and/or the traces for configured filters are obtained for the packet matching configured Source / Destination UDP Port.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

Interfaces

Use the Interfaces page to associate the CPU filters to interface or list of interfaces. The interfaces can be physical or logical LAG. The statistics counters are updated only for the configured interfaces. Similarly, the traces can also be obtained for configured interfaces.

To access this page, click **System > Advanced Configuration > CPU Traffic Filter > Interfaces**.




Figure 4.85 System > Advanced Configuration > CPU Traffic Filter > Interfaces
The following table describes the items in the previous figure.

Item	Description
Interface	The interfaces can be physical or logical LAG.
Direction	Only two software filter is supported (one filter each direction Tx or Rx) with condition matching as one, many or all in the below list in Tx or Rx or Both direction.
Remove	Click Remove to remove the selected entries.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new CPU Traffic filter to interface(s). See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

To add a new CPU Traffic filter to interface(s):

Click **System > Advanced Configuration > CPU Traffic Filter > Interfaces > Add**.

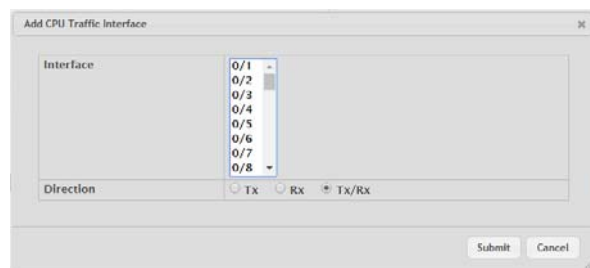


Figure 4.86 System > Advanced Configuration > CPU Traffic Filter > Interfaces > Add

The following table describes the items in the previous figure.

Item	Description
Interface	The interfaces can be physical or logical LAG.
Direction	Only two software filter is supported (one filter each direction Tx or Rx) with condition matching as one, many or all in the below list in Tx or Rx or Both direction.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

Statistics

Use the Statistics page to view per interface statistics for configured CPU filters.

To access this page, click **System > Advanced Configuration > CPU Traffic Filter > Statistics**.




Figure 4.87 System > Advanced Configuration > CPU Traffic Filter > Statistics

The following table describes the items in the previous figure.

Item	Description
Filter Name	The list of available filter names. Select a filter to view its interface statistics.
Interface	The interfaces can be physical or logical LAG.
Tx	The counter statistics for an interface associated with Tx direction.
Last Updated Tx Timestamp	Indicates the time when the sent packet count on an interface was last updated, based on the user defined packet filter on the interface.
Rx	The counter statistics for an interface associated with Rx direction.
Last Updated Rx Timestamp	Indicates the time when the received packet count on an interface was last updated, based on the user defined packet filter on the interface.
Refresh	Click Refresh to update the screen.
Clear	Click Clear to clear the interface statistics.

Summary

Use the Summary page to view all interface summary for CPU filters.

To access this page, click **System > Advanced Configuration > CPU Traffic Filter > Summary**.

LACPDU	0	0
ARP	0	0
UDLD	0	0
LLDP	0	0
IP	0	0
OSPF	0	0
BGP	0	0
DHCP	0	0
Broadcast	0	0

Figure 4.88 System > Advanced Configuration > CPU Traffic Filter > Summary

The following table describes the items in the previous figure.

Item	Description
Filter Name	The associated filter name.
Transmitted	The counter statistics for all interfaces, which are associated with Tx direction.
Received	The counter statistics for all interfaces, which are associated with Rx direction.
Refresh	Click Refresh to update the screen.
Clear	Click Clear to clear the filter summary.

Trace Information

Use the Trace Information page to view CPU Trace information.

To access this page, click **System > Advanced Configuration > CPU Traffic Filter > Trace Information**.



Figure 4.89 System > Advanced Configuration > CPU Traffic Filter > Trace Information

The following table describes the items in the previous figure.

Item	Description
Trace Information	It provides trace information for the matching packets as defined in the filters until the packet is delivered to registered application.
Refresh	Click Refresh to update the screen.
Clear	Click Clear to clear the CPU trace information.

4.3.3 Basic Configuration

4.3.3.1 Switch

IEEE 802.3x flow control works by pausing a port when the port becomes oversubscribed and dropping all traffic for small bursts of time during the congestion condition. This can lead to high-priority and/or network control traffic loss. When 802.3x flow control is enabled, lower speed switches can communicate with higher speed switches by requesting that the higher speed switch refrains from sending packets. Transmissions are temporarily halted to prevent buffer overflows.

To access this page, click **System > Basic Configuration > Switch**.




Figure 4.90 System > Basic Configuration > Switch

The following table describes the items in the previous figure.

Item	Description
802.3x Flow Control Mode	The 802.3x flow control mode on the switch. IEEE 802.3x flow control works by pausing a port when the port becomes oversubscribed. This allows lower-speed switches to communicate with higher-speed switches. A lower-speed or congested switch can send a PAUSE frame requesting that the peer device refrain from sending packets. Transmissions are temporarily halted to prevent buffer overflows. The options are as follows: <ul style="list-style-type: none">■ Disabled: The switch does not send PAUSE frames if the port buffers become full.■ Enabled: The switch can send PAUSE frames to a peer device if the port buffers become full.
MAC Address Aging Interval (Seconds)	The MAC address table (forwarding database) contains static entries, which never age out, and dynamically-learned entries, which are removed if they are not updated within a given time. Specify the number of seconds a dynamic address should remain in the MAC address table after it has been learned.

Item	Description
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.4 Configuration Storage

4.3.4.1 Save

Use the Save All Applied Changes page to store the system's configuration settings to non-volatile memory. Once saved the settings are available across a system reset. When you click **Save**, the save action is initiated.

To access this page, click **System > Configuration Storage > Save**.

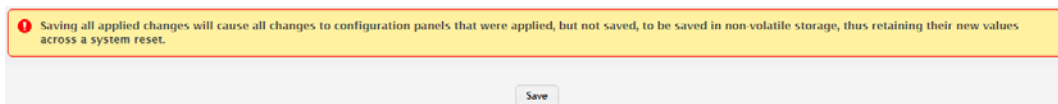


Figure 4.91 System > Configuration Storage > Save

The following table describes the items in the previous figure.

Item	Description
Save	Click Save to initiate a save of all system configuration after displaying a confirmation message. All of the current system configuration settings, including any that have been changed by the user, are stored into non-volatile memory so that they are preserved across a system reset.

4.3.4.2 Reset

Use the Reset Configuration page to reset the system's parameters to the factory default settings. The Reset function overrides all previously saved configuration changes. When you click **Reset**, the reset action is initiated.

To access this page, click **System > Configuration Storage > Reset**.



Figure 4.92 System > Configuration Storage > Reset

The following table describes the items in the previous figure.

Item	Description
Reset	Click Reset to initiate the action to reset all configuration parameters to their factory default settings after displaying a confirmation message. All configuration changes, including those that were previously saved, are reset in the running system by this action. It is possible that the IP address of the switch will change. If this occurs you will need to determine the new IP address to access the device using the web.

4.3.4.3 Erase Startup

Use the Erase Startup page to delete the text-based configuration file. The file is stored in non-volatile memory. When you click **Reset**, the Erase Startup action is initiated.

To access this page, click **System > Configuration Storage > Erase Startup**.




Figure 4.93 System > Configuration Storage > Erase Startup

The following table describes the items in the previous figure.

Item	Description
Reset	Click Reset to initiate the action to erase the text-based configuration file stored in non-volatile memory after displaying a confirmation message. If the system resets and no startup-config file is found, the system will begin the AutoInstall process to automatically update the image and download a configuration file.

4.3.4.4 Copy

Use the Copy Configuration Files page to copy the information contained in one configuration file to another configuration file on the device. When you click **Submit**, the copy action takes place immediately, and the source file overwrites the destination file.

To access this page, click **System > Configuration Storage > Copy**.




Figure 4.94 System > Configuration Storage > Copy

The following table describes the items in the previous figure.

Item	Description
Source File	Select the configuration file that will overwrite the contents in the selected destination file. The source file options are as follows: <ul style="list-style-type: none">■ Running Config: The file that contains the configuration that is currently active on the system. Copying the Running Config file to the Startup Config file is effectively the same as performing a Save.■ Startup Config: The file that contains the configuration that loads when the system boots.■ Backup Config: The file that is used to store a copy of the running or startup configuration.
Destination File	Select file to be overwritten by the contents in the selected source file. The destination file options are as follows: <ul style="list-style-type: none">■ Startup Config: The file that contains the configuration that loads when the system boots.■ Backup Config: The file that is used to store a copy of the running or startup configuration.
Submit	Click Submit to save the values and update the screen.

4.3.5 Connectivity

4.3.5.1 IPv4

Use the IPv4 Network Connectivity page to configure and view the IPv4 network connectivity information on the network interface. The network interface is the logical interface that allows remote management of the device via any of the front-panel switch ports. To enable management of the device over an IPv4 network by using a Web browser, SNMP, Telnet, or SSH, you must first configure it with an IP address, subnet mask, and default gateway. The configuration parameters associated with the network interface do not affect the configuration of the front-panel ports through which traffic is switched or routed.

To access this page, click **System > Connectivity > IPv4**.

Network Configuration Protocol	* None <input type="radio"/> Bootp <input type="radio"/> DHCP <input type="radio"/>	
DHCP Client Identifier	<input type="checkbox"/>	
IP Address	192.168.1.158	(x,x,x,x)
Subnet Mask	255.255.255.0	(x,x,x,x)
Default Gateway		(x,x,x,x)
MAC Address Type	* Burned In <input type="radio"/> Locally Administered <input type="radio"/>	
Burned In MAC Address	00:11:22:33:44:55	
Locally Administered MAC Address	00:00:00:00:00:00	(xxxxxxxxxxxx)(bit format of the first byte shall be 'xxxxxx10')
Management VLAN ID	1	(1 to 4093)

Figure 4.95 System > Connectivity > IPv4

The following table describes the items in the previous figure.

Item	Description
Network Configuration Protocol	Specify how the device acquires network information on the network interface: <ul style="list-style-type: none"> ■ None: The device does not attempt to acquire network information dynamically. Select this option to configure a static IP address, subnet mask, and default gateway. ■ BOOTP: During the next boot cycle, the BOOTP client on the device broadcasts a BOOTP request in an attempt to acquire information from a BOOTP server on the network. ■ DHCP: During the next boot cycle, the DHCP client on the device broadcasts a DHCP request in an attempt to acquire information from a DHCP server on the network. After this option is applied, you can use the Refresh icon at the end of the row to renew the IPv4 address learned from DHCP server.
DHCP Client Identifier	The DHCP Client Identifier (Option 61) is used by DHCP clients to specify their unique identifier. DHCP servers use this value to index their database of address bindings. This value is expected to be unique for all clients in an administrative domain. The Client Identifier string will be displayed beside the check box once DHCP is enabled on the port on which the Client Identifier option is selected. This web page will need to be refreshed once this change is made.
IP Address	The IP address of the interface. If the Network Configuration Protocol is None, you can manually configure a static IP address. If the Network Configuration Protocol is BOOTP or DHCP, this field displays the IP address that was dynamically acquired (if any).
Subnet Mask	The IP subnet mask for the interface. If the Network Configuration Protocol is None, you can manually configure a static subnet mask. If the Network Configuration Protocol is BOOTP or DHCP, this field displays the subnet mask that was dynamically acquired (if any).

Item	Description
Default Gateway	The default gateway for the IP interface. If the Network Configuration Protocol is None, you can manually configure the IP address of the default gateway. If the Network Configuration Protocol is BOOTP or DHCP, this field displays the default gateway address that was dynamically acquired (if any).
MAC Address Type	Specify whether the burned in or the locally administered MAC address should be used for in-band connectivity.
Burned In MAC Address	The burned in MAC address used for in-band connectivity if you choose not to configure a locally administered address.
Locally Administered MAC Address	You may configure a locally administered MAC address for in-band connectivity instead of using the burned in universally administered MAC address. In addition to entering an address in this field, you must also set the MAC address type to locally administered. Enter the address as twelve hexadecimal digits (6 bytes) with a colon between each byte. Bit 6 of byte 0 must be set to 1 and bit 0 to 0, i.e. byte 0 must have a value of 2, 6, A or E for its second digit.
Management VLAN ID	The VLAN ID for the management VLAN. Some network administrators use a management VLAN to isolate system management traffic from end-user data traffic.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.5.2 IPv6






Use the IPv6 Network Connectivity page to configure and view IPv6 information on the network interface. The network interface is the logical interface that allows remote management of the device via any of the front-panel switch ports. To enable management of the device over an IPv6 network by using a Web browser, SNMP, Telnet, or SSH, you must first configure the device with the appropriate IPv6 information. The configuration parameters associated with the network interface do not affect the configuration of the front-panel ports through which traffic is switched or routed.

To access this page, click **System > Connectivity > IPv6**.

Figure 4.96 System > Connectivity > IPv6

The following table describes the items in the previous figure.

Item	Description
IPv6 Mode	Enables or disables the IPv6 administrative mode on the network interface.
Network Configuration Protocol	Specify whether the device should attempt to acquire network information from a DHCPv6 server. Selecting None disables the DHCPv6 client on the network interface.

Item	Description
IPv6 Stateless Address AutoConfig Mode	<p>Sets the IPv6 stateless address auto configuration mode on the network interface.</p> <ul style="list-style-type: none"> Enabled: The network interface can acquire an IPv6 address through IPv6 Neighbor Discovery Protocol (NDP) and through the use of Router Advertisement messages. Disabled: The network interface will not use the native IPv6 address auto configuration features to acquire an IPv6 address.
DHCPv6 Client DUID	The client identifier used by the DHCPv6 client (if enabled) when sending messages to the DHCPv6 server.
IPv6 Gateway	The default gateway for the IPv6 network interface. To configure this field, click  button in the row. To reset the field to the default value, click  button in the row.
Static IPv6 Addresses	<p>Lists the manually configured static IPv6 addresses on the network interface.</p> <ul style="list-style-type: none"> To add an entry to the list, click  button to open the Add IPv6 Address dialog and provide the following: <ul style="list-style-type: none"> New IPv6 Address: Specify the IPv6 address to add to the interface. EUI Flag: Select this option to enable the Extended Universal Identifier (EUI) flag for IPv6 address, or clear the option to omit the flag. To delete an entry from the list, click  button associated with the entry to remove. To delete all entries from the list, click  button in the heading row.
Dynamic IPv6 Addresses	Lists the IPv6 addresses on the network interface that have been dynamically configured through IPv6 auto configuration or DHCPv6.
Default IPv6 Routers	Lists the IPv6 address of each default router that has been automatically configured through IPv6 router discovery.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.5.3 IPv6 Neighbors

When IPv6 is enabled on the service port, and a ping is initiated to a neighbor, the neighbor is added to the cache (if successful). The Network Port IPv6 Neighbors page displays data on these ports.

To access this page, click **System > Connectivity > IPv6 Neighbors**.




Figure 4.97 System > Connectivity > IPv6 Neighbors

The following table describes the items in the previous figure.

Item	Description
IPv6 Address	The IPv6 address of a neighbor device that has been reachable on the local link through the network interface.

Item	Description
MAC Address	The MAC address of the neighboring device.
Type	The type of the neighbor entry, which is one of the following: <ul style="list-style-type: none"> ■ Static: The neighbor entry is manually configured. ■ Dynamic: The neighbor entry is dynamically resolved. ■ Local: The neighbor entry is a local entry. ■ Other: The neighbor entry is an unknown entry.
Is Router	Identifies whether the neighbor device is a router. The possible values are: <ul style="list-style-type: none"> ■ True: The neighbor device is a router. ■ False: The neighbor device is not a router.
Neighbor State	The current reachability state of the neighboring device, which is one of the following: <ul style="list-style-type: none"> ■ Reachable: The neighbor is reachable through the network interface. ■ Stale: The neighbor is not known to be reachable, and the system will begin the process to reach the neighbor. ■ Delay: The neighbor is not known to be reachable, and upper-layer protocols are attempting to provide reachability information. ■ Probe: The neighbor is not known to be reachable, and the device is attempting to probe for this neighbor. ■ Unknown: The reachability status cannot be determined.
Last Updated	The amount of time that has passed since the neighbor entry was last updated.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new network port IPv6 neighbor. See the following procedure.
Remove	Click Remove to remove the selected entries.

To add a new network port IPv6 neighbor:

Click **System > Connectivity > IPv6 Neighbors > Add**.

Figure 4.98 System > Connectivity > IPv6 Neighbors > Add

The following table describes the items in the previous figure.

Item	Description
IPv6 Address	The IPv6 address of a neighbor device that has been reachable on the local link through the network interface.
MAC Address	The MAC address of the neighboring device.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.3.5.4 Service Port IPv4

Use the Service Port IPv4 Configuration page to configure network information on the service port. The service port is a dedicated Ethernet port for out-of-band management of the device. Traffic on this port is segregated from operational network traffic on the switch ports and cannot be switched or routed to the operational network.

To access this page, click **System > Connectivity > Service Port IPv4**.

Figure 4.99 System > Connectivity > Service Port IPv4

The following table describes the items in the previous figure.

Item	Description
Service Port Configuration Protocol	Specify how the device acquires network information on the service port: <ul style="list-style-type: none"> ■ BOOTP: During the next boot cycle, the BOOTP client on the device broadcasts a BOOTP request in an attempt to acquire information from a BootP server on the network. ■ DHCP: During the next boot cycle, the DHCP client on the device broadcasts a DHCP request in an attempt to acquire information from a DHCP server on the network. ■ None: The device does not attempt to acquire network information dynamically.
DHCP Client Identifier	The DHCP Client Identifier (Option 61) is used by DHCP clients to specify their unique identifier. DHCP servers use this value to index their database of address bindings. This value is expected to be unique for all clients in an administrative domain. The Client Identifier string will be displayed beside the check box once DHCP is enabled on the port on which the Client Identifier option is selected. This web page will need to be refreshed once this change is made.
IP Address	The IP address of the interface. If the Service Port Configuration Protocol is None, you can manually configure a static IP address. If the Service Port Configuration Protocol is BOOTP or DHCP, this field displays the IP address that was dynamically acquired (if any).
Subnet Mask	The IP subnet mask for the interface. If the Service Port Configuration Protocol is None, you can manually configure a static subnet mask. If the Service Port Configuration Protocol is BOOTP or DHCP, this field displays the subnet mask that was dynamically acquired (if any).
Default Gateway	The default gateway for the IP interface. If the Service Port Configuration Protocol is None, you can manually configure the IP address of the default gateway. If the Service Port Configuration Protocol is BOOTP or DHCP, this field displays the default gateway address that was dynamically acquired (if any).
Interface Status	Indicates whether the link status is up or down.
Burned In MAC Address	The burned in MAC address used for out-of-band connectivity.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Renew DHCP Lease	Click Renew DHCP Lease to renew DHCP lease.
Cancel	Click Cancel to restore default value.






4.3.5.5 Service Port IPv6

Use the Service Port IPv6 Configuration page to configure IPv6 network information on the service port. The service port is a dedicated Ethernet port for out-of-band management of the device. Traffic on this port is segregated from operational network traffic on the switch ports and cannot be switched or routed to the operational network.

To access this page, click **System > Connectivity > Service Port IPv6**.

Figure 4.100 System > Connectivity > Service Port IPv6

The following table describes the items in the previous figure.

Item	Description
IPv6 Mode	Enables or disables the IPv6 administrative mode on the service port.
Service Port Configuration Protocol	Specify whether the device should attempt to acquire network information from a DHCPv6 server. Selecting None disables the DHCPv6 client on the service port.
IPv6 Stateless Address AutoConfig Mode	<p>Sets the IPv6 stateless address auto configuration mode on the service port.</p> <ul style="list-style-type: none"> Enabled: The service port can acquire an IPv6 address through IPv6 Neighbor Discovery Protocol (NDP) and through the use of Router Advertisement messages. Disabled: The service port will not use the native IPv6 address auto configuration features to acquire an IPv6 address.
DHCPv6 Client DUID	The client identifier used by the DHCPv6 client (if enabled) when sending messages to the DHCPv6 server.
IPv6 Gateway	The default gateway for the IPv6 service port interface. To configure this field, click  button in the row. To reset the field to the default value, click  button in the row.
Static IPv6 Addresses	<p>Lists the manually configured static IPv6 addresses on the service port interface.</p> <ul style="list-style-type: none"> To add an entry to the list, click  button to open the Add IPv6 Address dialog and provide the following: <ul style="list-style-type: none"> New IPv6 Address: Specify the IPv6 address to add to the interface. EUI Flag: Select this option to enable the Extended Universal Identifier (EUI) flag for IPv6 address, or clear the option to omit the flag. To delete an entry from the list, click  button associated with the entry to remove. To delete all entries from the list, click  button in the heading row.
Dynamic IPv6 Addresses	Lists the IPv6 addresses on the service port interface that have been dynamically configured through IPv6 auto configuration or DHCPv6.
Default IPv6 Routers	Lists the IPv6 address of each default router that has been automatically configured through IPv6 router discovery.

Item	Description
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.5.6 Service Port IPv6 Neighbors

The Service Port IPv6 Neighbors page provides information about IPv6 neighbors the device has discovered through the service port by using the Neighbor Discovery Protocol (NDP). The manually configured static service port IPv6 neighbors are also displayed.

To access this page, click **System > Connectivity > Service Port IPv6 Neighbors**.




Figure 4.101 System > Connectivity > Service Port IPv6 Neighbors

The following table describes the items in the previous figure.

Item	Description
IPv6 Addresses	The IPv6 address of a neighbor device that has been reachable on the local link through the service port.
MAC Address	The MAC address of the neighboring device.
Type	The type of the neighbor entry, which is one of the following: <ul style="list-style-type: none"> ■ Static: The neighbor entry is manually configured. ■ Dynamic: The neighbor entry is dynamically resolved. ■ Local: The neighbor entry is a local entry. ■ Other: The neighbor entry is an unknown entry.
Is Router	Identifies whether the neighbor device is a router. The possible values are: <ul style="list-style-type: none"> ■ True: The neighbor device is a router. ■ False: The neighbor device is not a router.
Neighbor State	The current reachability state of the neighboring device, which is one of the following: <ul style="list-style-type: none"> ■ Reachable: The neighbor is reachable through the service port. ■ Stale: The neighbor is not known to be reachable, and the system will begin the process to reach the neighbor. ■ Delay: The neighbor is not known to be reachable, and upper-layer protocols are attempting to provide reachability information. ■ Probe: The neighbor is not known to be reachable, and the device is attempting to probe for this neighbor. ■ Unknown: The reachability status cannot be determined.
Last Updated	The amount of time that has passed since the neighbor entry was last updated.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new service port IPv6 neighbor. See the following procedure.
Remove	Click Remove to remove the selected entries.

To add a new service port IPv6 neighbor:

Click **System > Connectivity > Service Port IPv6 Neighbors > Add**.

Figure 4.102 System > Connectivity > Service Port IPv6 Neighbors List > Add

The following table describes the items in the previous figure.

Item	Description
IPv6 Address	The IPv6 address of a neighbor device that has been reachable on the local link through the service port.
MAC Address	The MAC address of the neighboring device.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.3.5.7 DHCP Client Options

Use the DHCP Client Options page to configure DHCP client settings on the system.

To access this page, click **System > Connectivity > DHCP Client Options**.

Figure 4.103 System > Connectivity > DHCP Client Options

The following table describes the items in the previous figure.

Item	Description
DHCP Vendor Class ID Mode	The VCI administrative mode. When the mode is enabled, the DHCP client includes the text configured as the DHCP Vendor Class ID String in DHCP requests.
DHCP Vendor Class ID String	The text string to add to DHCP requests as option 60, the VCI option.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.6 Firmware

4.3.6.1 Status

Use the Dual Image Status page to view information about the software images on the device. The device can store up to two software images in permanent storage. The dual image feature allows you to upgrade the device without deleting the older software image.

To access this page, click **System > Firmware > Status**.

Active	Backup	Current Active	Next Active
2.00.005	2.00.005	2.00.005	2.00.005
Image Description			
Active			
Backup			

Refresh

Figure 4.104 System > Firmware > Status

The following table describes the items in the previous figure.

Item	Description
Unit	The unit ID of the switch.
Active	The code file version of the active image.
Backup	The code file version of the backup image.
Current Active	The image version that is loaded and running on this unit.
Next Active	The image version to be loaded after the system reboots.
Image Description	
Active	The description associated with the active code file.
Backup	The description associated with the backup code file.
Refresh	Click Refresh to update the screen.

4.3.6.2 Configuration and Upgrade

Use the Dual Image Configuration and Upgrade page to transfer a new firmware (code) image to the device, select which image to load during the next boot cycle, and add a description to each image on the device. The device uses the HTTP protocol to transfer the image, and the image is saved as the backup image.

To access this page, click **System > Firmware > Configuration and Upgrade**.

Figure 4.105 System > Firmware > Configuration and Upgrade

The following table describes the items in the previous figure.

Item	Description
Images	
Active	<p>The active code file version. Use the icons to the right of the field to perform the file transfer.</p> <ul style="list-style-type: none"> To transfer a new code image to the device, click button. The Firmware Upgrade window opens. Click Choose File to browse to the file to transfer. After you select the appropriate file, click Begin Transfer to launch the HTTP transfer process. The active image is overwritten by the file that you transfer.
Backup	<p>The backup code file version. Use the icons to the right of the field to perform the following tasks:</p> <ul style="list-style-type: none"> To transfer a new code image to the device, click button. The Firmware Upgrade window opens. Click Choose File to browse to the file to transfer. After you select the appropriate file, click Begin Transfer to launch the HTTP transfer process. If a backup image already exists on the device, it is overwritten by the file that you transfer. To delete the backup image from permanent storage, click button. You must confirm the action before the image is deleted.
Next Active	Select the image version to load the next time this unit reboots.

Item	Description
Image Description	
Active	Specify a description to associate with the image that is currently the active code file.
Backup	Specify a description to associate with the image that is currently the backup code file.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.6.3 AutoInstall

The AutoInstall feature can automatically obtain configuration information and install a new image when the switch boots. The process begins when the switch is initialized and no configuration file (startup-config) is found, or when the switch boots and loads a saved configuration that has AutoInstall enabled. If initiated, the AutoInstall feature allows the device to obtain an IP address from a network DHCP server and then attempts to locate the predefined configuration file from a TFTP server.

To access this page, click **System > Firmware > AutoInstall**.

Figure 4.106 System > Firmware > AutoInstall

The following table describes the items in the previous figure.

Item	Description
Admin Mode	The current administrative mode of the AutoInstall feature: <ul style="list-style-type: none"> Start: AutoInstall is enabled, and the feature will attempt to automatically configure the device during the next boot cycle. Stop: AutoInstall is disabled. The automatic process will begin only if no configuration file is located during the next boot cycle.
Persistent Mode	If this option is selected, the settings you configure on this page are automatically saved to persistent memory in the startup-config file when you apply the changes. If this option is not selected, the device treats these settings like any other applied changes (i.e. the changes are not retained across a reboot unless you save the configuration).
AutoSave Mode	If this option is selected, the downloaded configuration is automatically saved to persistent storage. If this option is not selected, you must explicitly save the downloaded configuration in non-volatile memory for the configuration to be available for the next reboot.
AutoReboot Mode	If this option is selected, the switch automatically reboots after a new image is successfully downloaded and makes the downloaded image the active image. If this option is not selected, the device continues to boot with the current image. The downloaded image will not become the active image until the device reboots.
Retry Count	When attempting to retrieve the DHCP-specified configuration file, this value represents the number of times the TFTP client on the device tries to use unicast requests before reverting to broadcast requests.
Status	The current status of the AutoInstall process.
Submit	Click Submit to save the values and update the screen.

Item	Description
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.7 Logs

4.3.7.1 Buffered Log

The log messages the device generates in response to events, faults, errors, and configuration changes are stored locally on the device in the RAM (cache). This collection of log files is called the RAM log or buffered log. When the buffered log file reaches the configured maximum size, the oldest message is deleted from the RAM when a new message is added. If the system restarts, all messages are cleared.

To access this page, click **System > Logs > Buffered Log**.

Log Index	Log Time	Severity	Component	Description
1	Jan 1 18:08:54	Info	General	Application Started (opensshd, ID = 9, PID = 799)
2	Jan 1 18:08:54	Info	General	Application Terminated (opensshd, ID = 9, PID = 798)
3	Jan 1 18:08:53	Info	General	Application Started (opensshd, ID = 9, PID = 798)
4	Jan 1 18:08:53	Info	General	Application Terminated (opensshd, ID = 9, PID = 797)
5	Jan 1 18:08:53	Info	General	Application Started (opensshd, ID = 9, PID = 797)
6	Jan 1 18:08:53	Info	General	Application Terminated (opensshd, ID = 9, PID = 796)
7	Jan 1 18:08:52	Info	General	Application Started (opensshd, ID = 9, PID = 796)
8	Jan 1 18:08:52	Info	General	Application Terminated (opensshd, ID = 9, PID = 795)
9	Jan 1 18:08:52	Info	General	Application Started (opensshd, ID = 9, PID = 795)
10	Jan 1 18:08:52	Info	General	Application Terminated (opensshd, ID = 9, PID = 794)

Figure 4.107 System > Logs > Buffered Log

The following table describes the items in the previous figure.

Item	Description
Log Index	The position of the entry within the buffered log file. The most recent log message always has a Log Index value of 1.
Log Time	The time the entry was added to the log.
Severity	The severity level associated with the log entry. The severity can be one of the following: <ul style="list-style-type: none"> ■ Emergency (0): The device is unusable. ■ Alert (1): Action must be taken immediately. ■ Critical (2): The device is experiencing primary system failures. ■ Error (3): The device is experiencing non-urgent failures. ■ Warning (4): The device is experiencing conditions that could lead to system errors if no action is taken. ■ Notice (5): The device is experiencing normal but significant conditions. ■ Info (6): The device is providing non-critical information. ■ Debug (7): The device is providing debug-level information.
Component	The component that issued the log entry.
Description	The text description for the log entry.
Refresh	Click Refresh to update the screen.
Clear Log	Click Clear Log to clear the buffered log messages and resets the counters.

4.3.7.2 Event Log

Use the Event Log page to display the event log, which is used to hold error messages for catastrophic events. After the event is logged and the updated log is saved in flash memory, the switch will be reset. The log can hold at least 2,000 entries (the actual number depends on the platform and OS), and is erased when an attempt is made to add an entry after it is full. The event log is preserved across system resets.

To access this page, click **System > Logs > Event Log**.

Log Index	Type	Filename	Line	Task ID	Code	Event Time
1	EVENT	usmdb_sim.c	3632	049B4D24	00000000	0d:23:16:01
2	EVENT	usmdb_sim.c	3632	03544D24	00000000	0d:18:21:11
3	EVENT	usmdb_sim.c	3632	050A5D24	00000000	3d:22:22:54
4	EVENT	usmdb_sim.c	3632	03838D24	00000000	0d:01:27:23
5	EVENT	usmdb_sim.c	3632	047FED24	00000000	0d:00:25:28
6	EVENT	usmdb_sim.c	3632	04766D24	00000000	0d:19:46:05
7	EVENT	usmdb_sim.c	3632	04D94D34	00000000	0d:00:06:48
8	EVENT	usmdb_sim.c	3632	04986D24	00000000	0d:00:14:13
9	EVENT	usmdb_sim.c	3632	04162D34	00000000	0d:05:18:27
10	EVENT	usmdb_sim.c	3632	035E1D34	00000000	0d:00:18:11

Figure 4.108 System > Logs > Event Log

The following table describes the items in the previous figure.

Item	Description
Log Index	A display row index number used to identify the event log entry, with the most recent entry listed first (lowest number).
Type	The incident category that indicates the cause of the log entry: EVENT, ERROR, etc.
Filename	The source code filename of the event origin.
Line	Within the source code filename, the line number of the event origin.
Task ID	A system identifier of the task that was running when the event occurred. This value is assigned by, and is specific to, the operating system.
Code	An event-specific code value that is passed to the log handler by the source code file reporting the event.
Event Time	A time stamp (days, hours, minutes, and seconds) indicating when the event occurred, measured from the time the device was last reset. The only correlation between any two entries in the event log is the relative amount of time after a system reset that the event occurred.
Refresh	Click Refresh to update the screen.

4.3.7.3 Persistent Log

Persistent log messages are stored in persistent storage so that they survive across device reboots. Two types of log files exist in flash (persistent) memory: the system startup log and the system operation logs. The system startup log stores the first 32 messages received after system reboot. The log file stops when it is full. The system operation log stores the last 32 messages received during system operation. The oldest messages are overwritten when the file is full.

To access this page, click **System > Logs > Persistent Log**.




Figure 4.109 System > Logs > Persistent Log

The following table describes the items in the previous figure.

Item	Description
Log Index	The position of the entry within the buffered log file. The most recent log message always has a Log Index value of 1.
Log Time	The time the entry was added to the log.
Severity	The severity level associated with the log entry. The severity can be one of the following: <ul style="list-style-type: none"> ■ Emergency (0): The device is unusable. ■ Alert (1): Action must be taken immediately. ■ Critical (2): The device is experiencing primary system failures. ■ Error (3): The device is experiencing non-urgent failures. ■ Warning (4): The device is experiencing conditions that could lead to system errors if no action is taken. ■ Notice (5): The device is experiencing normal but significant conditions. ■ Info (6): The device is providing non-critical information. ■ Debug (7): The device is providing debug-level information.
Component	The component that has issued the log entry.
Description	The text description for the log entry.
Refresh	Click Refresh to update the screen.

4.3.7.4 Hosts

Use the Logging Hosts page to configure remote logging hosts where the switch can send logs.

To access this page, click **System > Logs > Hosts**.




Figure 4.110 System > Logs > Hosts

The following table describes the items in the previous figure.

Item	Description
Host	The IP address or DNS-resolvable host name of the remote host to receive log messages.

Item	Description
Status	Indicates whether the host has been configured to be actively logging or not.
Port	The UDP port on the logging host to which syslog messages are sent.
Severity Filter	Severity level threshold for log messages. All log messages with a severity level at and above the configured level are forwarded to the logging host.
Transport Mode	Transport mode used while sending messages to syslog servers. Supported modes are : UDP and TLS. If TLS is not configured default transport mode is UDP.
Authentication Mode	Using TLS security user can configure anonymous authentication mode, in which no client authentication is done by the syslog server. For x509/name authentication mode, two way authentication is done both by syslog client and client authentication by syslog server side.
Certificate Index	The index used for identifying corresponding certificate files.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new host. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

To add a new host:

Click **System > Logs > Hosts > Add**.

Figure 4.111 System > Logs > Hosts > Add

The following table describes the items in the previous figure.

Item	Description
IP Address/Host Name	The IP address or DNS-resolvable host name of the remote host to receive log messages.
Port	The UDP port on the logging host to which syslog messages are sent.
Severity Filter	Severity level threshold for log messages. All log messages with a severity level at and above the configured level are forwarded to the logging host.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.3.7.5 Configuration

The Log Configuration page allows administrators with the appropriate privilege level to configure the administrative mode and various settings for logging features on the switch.

To access this page, click **System > Logs > Configuration**.

The screenshot shows a web interface for configuring logging. It is divided into five sections:

- Buffered Log Configuration:** Admin Mode (radio buttons for Disable and Enable, with Enable selected), Behavior (radio buttons for Wrap and Stop on Full, with Wrap selected).
- Command Logger Configuration:** Admin Mode (radio buttons for Disable and Enable, with Disable selected).
- Console Log Configuration:** Admin Mode (radio buttons for Disable and Enable, with Enable selected), Severity Filter (dropdown menu set to Error).
- Persistent Log Configuration:** Admin Mode (radio buttons for Disable and Enable, with Disable selected), Severity Filter (dropdown menu set to Alert).
- Syslog Configuration:** Admin Mode (radio buttons for Disable and Enable, with Disable selected), Protocol Version (radio buttons for RFC 3164 and RFC 5424, with RFC 3164 selected), Local UDP Port (text input field with 514 and a note (1 to 65535)).

At the bottom, there are buttons for Submit, Refresh, and Cancel.

Figure 4.112 System > Logs > Configuration

The following table describes the items in the previous figure.

Item	Description
Buffered Log Configuration	
Admin Mode	Enable or disable logging to the buffered (RAM) log file.
Behavior	Specify what the device should do when the buffered log is full. It can either overwrite the oldest messages (Wrap) or stop writing new messages to the buffer (Stop on Full).
Command Logger Configuration	
Admin Mode	Enable or disable logging of the command-line interface (CLI) commands issued on the device.
Console Log Configuration	
Admin Mode	Enable or disable logging to any serial device attached to the host.
Severity Filter	Select the severity of the messages to be logged. All messages at and above the selected threshold are logged to the console. The severity can be one of the following: <ul style="list-style-type: none"> ■ Emergency (0): The device is unusable. ■ Alert (1): Action must be taken immediately. ■ Critical (2): The device is experiencing primary system failures. ■ Error (3): The device is experiencing non-urgent failures. ■ Warning (4): The device is experiencing conditions that could lead to system errors if no action is taken. ■ Notice (5): The device is experiencing normal but significant conditions. ■ Info (6): The device is providing non-critical information. ■ Debug (7): The device is providing debug-level information.
Persistent Log Configuration	
Admin Mode	Enable or disable logging to the persistent log. These messages are not deleted when the device reboots.
Severity Filter	Select the severity of the messages to be logged. All messages at and above the selected threshold are logged to the console. See the previous severity filter description for more information about each severity level.

Item	Description
Syslog Configuration	
Admin Mode	Enable or disable logging to configured syslog hosts. When the syslog admin mode is disabled the device does not relay logs to syslog hosts, and no messages will be sent to any collector/relay. When the syslog admin mode is enabled, messages will be sent to configured collectors/relays using the values configured for each collector/relay.
Protocol Version	The RFC version of the syslog protocol.
Local UDP Port	The UDP port on the local host from which syslog messages are sent.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.7.6 Source Interface Configuration

Use the Syslog Source Interface Configuration page to specify the physical or logical interface to use as the logging (Syslog) client source interface. When an IP address is configured on the source interface, this address is used for all Syslog communications between the local logging client and the remote Syslog server. The IP address of the designated source interface is used in the IP header of Syslog management protocol packets. This allows security devices, such as firewalls, to identify all source packets coming from a specific device.

To access this page, click **System > Logs > Source Interface Configuration**.

Figure 4.113 System > Logs > Source Interface Configuration

The following table describes the items in the previous figure.

Item	Description
Type	The type of interface to use as the source interface: <ul style="list-style-type: none"> ■ None: The primary IP address of the originating (outbound) interface is used as the source address. ■ Interface: The primary IP address of a physical port is used as the source address. ■ VLAN: The primary IP address of a VLAN routing interface is used as the source address. ■ Loopback: The primary IP address of the loopback interface is used as the source address. A loopback is always reachable, as long as any routing interface is up. ■ Tunnel: The primary IP address of a tunnel interface is used as the source address. ■ Network: The network source IP is used as the source address. ■ Service Port: The management port source IP is used as the source address.
Interface	When the selected Type is Interface, select the physical port to use as the source interface.
VLAN ID	When the selected Type is VLAN, select the VLAN to use as the source interface. The menu contains only the VLAN IDs for VLAN routing interfaces.

Item	Description
Loopback Interface	When the selected Type is Loopback, select the loopback interface to use as the source interface.
Tunnel ID	When the selected Type is Tunnel, select the tunnel interface to use as the source interface.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.7.7 Statistics

The Log Statistics page displays summary information about the number of messages logged to the buffered, persistent, or syslog file. It also displays the number of messages that were successfully or unsuccessfully relayed to any remote syslog servers configured on the device.

To access this page, click **System > Logs > Statistics**.

Buffered Log	
Total Number of Messages	42115
Persistent Log	
Total Number of Messages	0
Syslog	
Messages Received	42173
Messages Dropped	0
Messages Relayed	0
Refresh	

Figure 4.114 System > Logs > Statistics

The following table describes the items in the previous figure.

Item	Description
Buffered Log	
Total Number of Messages	The number of log messages currently stored in RAM.
Persistent Log	
Total Number of Messages	The number of log messages currently stored in persistent storage.
Syslog	
Messages Received	The total number of messages received by the log process. This includes messages that are dropped or ignored. The number includes messages of all severity levels.
Messages Dropped	The number of messages that failed to be relayed to a remote syslog server. The configured syslog server might be unreachable, misconfigured, or out of storage space.
Messages Relayed	The number of log messages successfully relayed to a remote syslog server. Messages forwarded to multiple hosts are counted once for each host.
Refresh	Click Refresh to update the screen.

4.3.8 Management Access

4.3.8.1 System

Use the System Connectivity page to control access to the management interface by administratively enabling or disabling various access methods.

To access this page, click **System > Management Access > System**.

The screenshot shows a configuration page with several sections, each with a title bar and a table of options. The sections are: HTTP (with 'HTTP Admin Mode' set to 'Enable'), Telnet (with 'Telnet Server Admin Mode' set to 'Disable' and 'Allow New Sessions' checked), Outbound Telnet (with 'Allow New Sessions' checked), Secure HTTP (with 'HTTPS Admin Mode' set to 'Disable'), and Secure Shell (with 'SSH Admin Mode' set to 'Enable'). At the bottom are 'Submit', 'Refresh', and 'Cancel' buttons.

Figure 4.115 System > Management Access > System

The following table describes the items in the previous figure.

Item	Description
HTTP	
HTTP Admin Mode	Enables or disables the HTTP administrative mode. When this mode is enabled, the device management interface can be accessed through a web browser using the HTTP protocol.
Telnet	
Telnet Server Admin Mode	Enables or disables the telnet administrative mode. When this mode is enabled, the device command-line interface (CLI) can be accessed through the telnet port. Disabling this mode disconnects all existing telnet connections and shuts down the telnet port in the device.
Allow New Sessions	Enables or disables new telnet sessions. When this option is disabled, the system does not accept any new telnet sessions, but existing telnet sessions are unaffected.
Outbound Telnet	
Allow New Sessions	Enables or disables new telnet sessions. When this option is disabled, the system does not accept any new telnet sessions, but existing telnet sessions are unaffected.
Secure HTTP	
HTTPS Admin Mode	Enables or disables the administrative mode of secure HTTP. When this mode is enabled, the device management interface can be accessed through a web browser using the HTTPS protocol.
Secure Shell	
SSH Admin Mode	Enables or disables the administrative mode of SSH. When this mode is disabled, all existing SSH connections remain connected until timed-out or logged out, but new SSH connections cannot be established.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.8.2 Telnet

Telnet is a terminal emulation TCP/IP protocol. ASCII terminals can be virtually connected to the local device through a TCP/IP protocol network. Telnet is an alternative to a local login terminal where a remote login is required.

The switch supports up to five simultaneous telnet sessions. All CLI commands can be used over a telnet session.

The Telnet Session Configuration page allows you to control inbound telnet settings on the switch. Inbound telnet sessions originate on a remote system and allow a user on that system to connect to the switch CLI.

To access this page, click **System > Management Access > Telnet**.

The screenshot shows a configuration form for Telnet settings. At the top, there is a radio button for 'Admin Mode' with 'Disable' selected and 'Enable' unselected. Below this are four rows of configuration fields: 'Telnet Port' with a value of 23 and a range of (1 to 65535, 23 = Default); 'Session Timeout (Minutes)' with a value of 5 and a range of (1 to 160); 'Maximum Number of Sessions' with a value of 5 and a range of (0 to 5); and 'Allow New Sessions' with a checked checkbox. At the bottom of the form are three buttons: 'Submit', 'Refresh', and 'Cancel'.

Figure 4.116 System > Management Access > Telnet

The following table describes the items in the previous figure.

Item	Description
Admin Mode	Enables or disables the telnet administrative mode. When enabled, the device may be accessed through the telnet port (23). Disabling this mode value disconnects all existing telnet connections and shuts down the telnet port in the device.
Telnet Port	The TCP port number on which the telnet server listens for requests. Existing telnet login sessions are not affected by a change in this value, although establishment of any new telnet sessions must use the new port number. <i>NOTE: Before changing this value, check your system (e.g. using netstat) to make sure the desired port number is not currently being used by any other service.</i>
Session Timeout (Minutes)	The telnet session inactivity timeout value, in minutes. A connected user that does not exhibit any telnet activity for this amount of time is automatically disconnected from the device.
Maximum Number of Sessions	The maximum number of telnet sessions that may be connected to the device simultaneously.
Allow New Sessions	Controls whether new telnet sessions are allowed. Setting this value to Disable disallows any new telnet sessions from starting (although existing telnet sessions are unaffected).
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.8.3 Outbound Telnet

The Outbound Telnet page displays the current value of the outbound Telnet settings on the device. An outbound Telnet session is a Telnet session initiated from the CLI of the device to the Telnet client on a remote device.

To access this page, click **System > Management Access > Outbound Telnet**.

The screenshot shows a configuration form for Outbound Telnet settings. It contains three rows of configuration fields: 'Allow New Sessions' with a checked checkbox; 'Maximum Number of Sessions' with a value of 5 and a range of (0 to 5); and 'Session Timeout (Minutes)' with a value of 5 and a range of (1 to 160). At the bottom of the form are three buttons: 'Submit', 'Refresh', and 'Cancel'.

Figure 4.117 System > Management Access > Outbound Telnet

The following table describes the items in the previous figure.

Item	Description
Allow New Sessions	Controls whether new outbound Telnet sessions are allowed. Setting this value to Disable disallows any new outbound Telnet sessions from starting (although existing Telnet sessions are unaffected).
Maximum Number of Sessions	The maximum number of allowed outbound Telnet sessions from the device simultaneously.
Session Timeout	Outbound telnet session inactivity timeout value, in minutes. An outbound Telnet session is closed automatically if there is no activity within the configured amount of time.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.8.4 Serial

The Serial Port page allows you to change the switch's serial port settings. In order for a terminal or terminal emulator to communicate with the switch, the serial port settings on both devices must be the same. Some settings on the switch cannot be changed.

To access this page, click **System > Management Access > Serial**.

The screenshot shows a configuration form for serial port settings. The fields and their values are: Serial Time Out (Minutes) is 5; Baud Rate (bps) is 115200; Character Size (Bits) is 8; Parity is None; Stop Bits is 1; and Flow Control is Disable. At the bottom of the form are three buttons: Submit, Refresh, and Cancel.

Figure 4.118 System > Management Access > Serial

The following table describes the items in the previous figure.

Item	Description
Serial Time Out (Minutes)	Serial port inactivity timeout value, in minutes. A logged-in user who does not exhibit any CLI activity through the serial port connection for this amount of time is automatically logged out of the device.
Baud Rate (bps)	The number of signals per second transmitted over the physical medium, measured in bits per second.
Character Size (Bits)	The number of bits in a character. This value is always 8.
Parity	The parity method used on the serial port.
Stop Bits	The number of stop bits per character.
Flow Control	Indicates whether hardware flow control is enabled or disabled on the serial port.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.8.5 CLI Banner

Use the CLI Banner Configuration page to configure a message that appears before the user prompt as a Pre- login banner. The message configured shows up on Telnet, SSH and Console connections.

To access this page, click **System > Management Access > CLI Banner**.

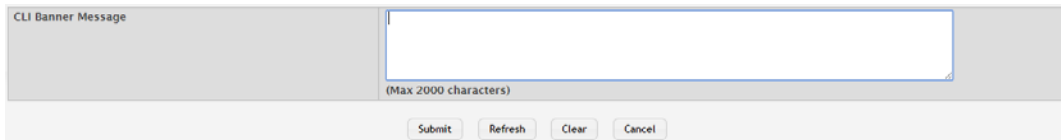


Figure 4.119 System > Management Access > CLI Banner

The following table describes the items in the previous figure.

Item	Description
CLI Banner Message	Text area for creating, viewing, or updating the CLI banner message. To create the CLI banner message, type the desired message in the text area. If you reach the end of the line, the text wraps to the next line. The line might not wrap at the same location in the CLI. To create a line break (carriage return) in the message, press the Enter key on the keyboard. The line break in the text area will be at the same location in the banner message when viewed through the CLI.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Clear	Click Clear to clear the CLI banner message from the device
Cancel	Click Cancel to restore default value.

4.3.8.6 HTTP

Use the HTTP Configuration page to configure the HTTP server settings on the system.

To access this page, click **System > Management Access > HTTP**.

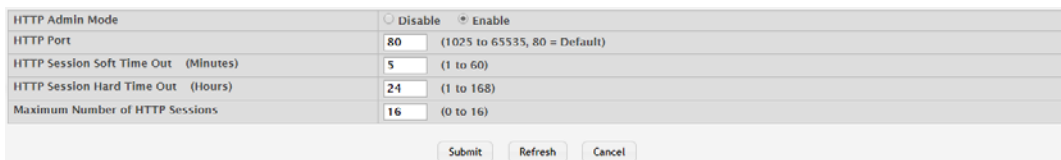


Figure 4.120 System > Management Access > HTTP

The following table describes the items in the previous figure.

Item	Description
HTTP Admin Mode	Enables or disables the HTTP administrative mode. When enabled, the device can be accessed through a web browser using the HTTP protocol.
HTTP Port	The TCP port number on which the HTTP server listens for requests. Existing HTTP login sessions are closed whenever this value is changed. All new HTTP sessions must use the new port number. <i>NOTE: Before changing this value, check your system (e.g. using netstat) to make sure the desired port number is not currently being used by any other service.</i>
HTTP Session Soft Time Out (Minutes)	HTTP session inactivity timeout value. A logged-in user that does not exhibit any HTTP activity for this amount of time is automatically logged out of the HTTP session.

Item	Description
HTTP Session Hard Time Out (Hours)	HTTP session hard timeout value. A user connected to the device via an HTTP session is automatically logged out after this amount of time regardless of the amount of HTTP activity that occurs.
Maximum Number of HTTP Sessions	The maximum number of HTTP sessions that may be connected to the device simultaneously.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.8.7 HTTPS

Use the Secure HTTP Configuration page to view and modify the Secure HTTP (HTTPS) settings on the device. HTTPS increases the security of web-based management by encrypting communication between the administrative system and the device.




To access this page, click **System > Management Access > HTTPS**.

HTTPS Admin Mode	<input checked="" type="radio"/> Disable <input type="radio"/> Enable
TLS Version 1	<input type="radio"/> Disable <input checked="" type="radio"/> Enable
SSL Version 3	<input type="radio"/> Disable <input checked="" type="radio"/> Enable
HTTPS Port	443 (1025 to 65535, 443 = Default)
HTTPS Session Soft Time Out (Minutes)	5 (1 to 60)
HTTPS Session Hard Time Out (Hours)	24 (1 to 168)
Maximum Number of HTTPS Sessions	5 (0 to 5)
Certificate Status	Absent

Figure 4.121 System > Management Access > HTTPS

The following table describes the items in the previous figure.

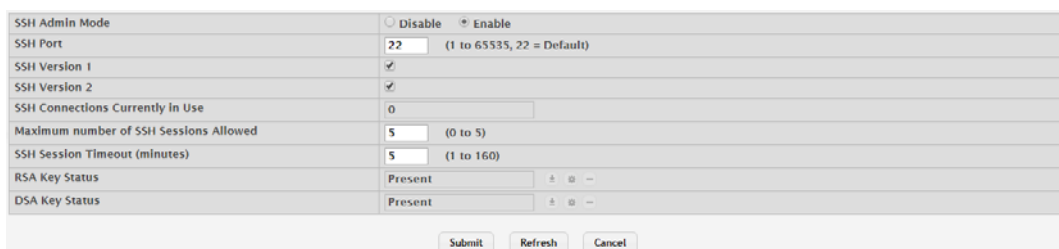
Item	Description
HTTPS Admin Mode	Enables or disables the HTTPS administrative mode. When this mode is enabled, the device can be accessed through a web browser using the HTTPS protocol.
TLS Version 1	Enables or disables Transport Layer Security Version 1.0. When this option is enabled, communication between the web browser on the administrative system and the web server on the device is sent through TLS 1.0.
SSL Version 3	Enables or disables Secure Sockets Layer Version 3.0. When this option is enabled, communication between the web browser on the administrative system and the web server on the device is sent through SSL 3.0. SSL must be administratively disabled while downloading an SSL certificate file from a remote server to the device.
HTTPS Port	The TCP port number that HTTPS uses. <i>NOTE: Before changing this value, check your system (e.g. using netstat) to make sure the desired port number is not currently being used by any other service.</i>
HTTPS Session Soft Time Out (Minutes)	HTTPS session inactivity timeout value. A logged-in user that does not exhibit any HTTPS activity for this amount of time is automatically logged out of the HTTPS session.
HTTPS Session Hard Time Out (Hours)	HTTPS session hard timeout value. A user connected to the device via an HTTPS session is automatically logged out after this amount of time regardless of the amount of HTTPS activity that occurs.
Maximum Number of HTTPS Sessions	The maximum number of HTTPS sessions that can be connected to the device simultaneously.

Item	Description
Certificate Status	The status of the SSL certificate generation process. <ul style="list-style-type: none"> ■ Present: The certificate has been generated and is present on the device. ■ Absent: Certificate is not available on the device. ■ Generation In Progress: An SSL certificate is currently being generated.
	Allows you to download an SSL certificate file from a remote system to the device. Note that to download SSL certificate files, SSL must be administratively disabled. <ul style="list-style-type: none"> ■ File Type: Specify the type of file to transfer from the device to a remote system. ■ Select File: Provides option to browse to the directory where the file is located and select the file to transfer to the device. ■ Status: Provides information about the status of the file transfer.
	Generates an SSL certificate to use for secure communication between the web browser and the embedded web server on the device.
	Deletes the SSL certificate. This button is available only if an SSL certificate is present on the device.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.8.8 SSH

Use the SSH Configuration page to view and modify the Secure Shell (SSH) server settings on the device. SSH is a network protocol that enables access to the CLI management interface by using an SSH client on a remote administrative system. SSH is a more secure access method than Telnet because it encrypts communication between the administrative system and the device. This page also allows you to download or generate SSH host keys for secure CLI-based management.

To access this page, click **System > Management Access > SSH**.



The screenshot shows the SSH Configuration page with the following settings:




- SSH Admin Mode: Disable Enable
- SSH Port: 22 (1 to 65535, 22 = Default)
- SSH Version 1:
- SSH Version 2:
- SSH Connections Currently In Use: 0
- Maximum number of SSH Sessions Allowed: 5 (0 to 5)
- SSH Session Timeout (minutes): 5 (1 to 160)
- RSA Key Status: Present
- DSA Key Status: Present

Buttons at the bottom: Submit, Refresh, Cancel.

Figure 4.122 System > Management Access > SSH

The following table describes the items in the previous figure.

Item	Description
SSH Admin Mode	Enables or disables the SSH server administrative mode. When this mode is enabled, the device can be accessed by using an SSH client on a remote system.

Item	Description
SSH Port	<p>The TCP port number on which the SSH server listens for requests. Existing SSH login sessions are not affected by a change in this value, although establishment of any new SSH sessions must use the new port number.</p> <p><i>NOTE: Before changing this value, check your system (e.g. using netstat) to make sure the desired port number is not currently being used by any other service.</i></p>
SSH Version 1	When this option is selected, the SSH server on the device can accept connections from an SSH client using SSH-1 protocol. If the option is clear, the device does not allow connections from clients using the SSH-1 protocol.
SSH Version 2	When this option is selected, the SSH server on the device can accept connections from an SSH client using SSH-2 protocol. If the option is clear, the device does not allow connections from clients using the SSH-2 protocol.
SSH Connections Currently in Use	The number of active SSH sessions between remote SSH clients and the SSH server on the device.
Maximum number of SSH Sessions Allowed	The maximum number of SSH sessions that may be connected to the device simultaneously.
SSH Session Timeout (minutes)	The SSH session inactivity timeout value. A connected user that does not exhibit any SSH activity for this amount of time is automatically disconnected from the device.
RSA Key Status	The status of the SSH-1 Rivest-Shamir-Adleman (RSA) key file or SSH-2 RSA key file (PEM Encoded) on the device, which might be Present, Absent, or Generation in Progress.
DSA Key Status	The status of the SSH-2 Digital Signature Algorithm (DSA) key file (PEM Encoded) on the device, which might be Present, Absent, or Generation in Progress.
	<p>Click the button to download an SSH-1 RSA, SSH-2 RSA, or SSH-2 DSA key file from a remote system to the device. After you click the button, a Download Certificates window opens. Select the file type to download, browse to the location on the remote system, and select the file to upload. Then, click Begin Transfer. The Status field provides information about the file transfer.</p> <ul style="list-style-type: none"> ■ File Type: Specify the type of file to transfer from the device to a remote system. ■ Select File: Provides option to browse to the directory where the file is located and select the file to transfer to the device. ■ Status: Provides information about the status of the file transfer.
	Click the button to delete an RSA key or DSA key that has been downloaded to the device or manually generated on the device.
	Deletes the SSL certificate. This button is available only if an SSL certificate is present on the device.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.9 Management Security

4.3.9.1 Access Profile

The administrator may elect to configure a management access control list. The Management Access Control and Administration List (ACAL) component is used to ensure that only known and trusted devices are allowed to remotely manage the switch via TCP/IP.

Note! Profile rules cannot be added or modified when a profile is active.



To access this page, click **System > Management Security > Access Profile**.

Figure 4.123 System > Management Security > Access Profile

The following table describes the items in the previous figure.

Item	Description
Profile Configuration	
Access Profile	Profile name for the Management Access Control list. One user defined Access Profile can be created.
Active Profile	Currently enabled profile name.
Packets Filtered	The number of packets filtered.
Profile Rule Configuration	
Interface	The port/interface or trunk ID.
Management Method	Below are the types of action will be taken on access control list. <ul style="list-style-type: none">■ Permit: To allow conditions for the management access list.■ Deny: To deny conditions for the management access list.
Source IP Address	IP Address of device which needs to permit or deny in the management access list.
Subnet Mask	Specifies the network mask of the source IP address.
VLAN	Vlan number.
Port Channel	Port channels, also known as Link Aggregation Groups (LAGs), allow one or more full-duplex Ethernet links of the same speed to be aggregated together.

Item	Description
Service	Indicates service type. Can be one of the following. <ul style="list-style-type: none"> ■ ANY ■ TELNET ■ HTTP ■ HTTPS ■ SNMP ■ SSH ■ TFTP ■ SNTP ■ JAVA
Priority	Priority for the rule. Duplicates are not allowed.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Add	Click Add to add Profile Rules to the Management Access Control list.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

4.3.10 Passwords

4.3.10.1 Line Password

Use the Line Password Configuration page to configure line mode passwords.

To access this page, click **System > Passwords > Line Password**.

Figure 4.124 System > Passwords > Line Password

The following table describes the items in the previous figure.

Item	Description
Line Mode	Any or all of the following passwords may be changed on this page by checking the box that precedes it: <ul style="list-style-type: none"> ■ Console ■ Telnet ■ SSH
Password	Enter the new password for the corresponding Line Mode in this field. Be sure the password conforms to the allowed number of characters. The password characters are not displayed on the page, but are disguised in a browser-specific manner.
Confirm Password	Re-enter the new password for the corresponding Line Mode in this field. This must be the same value entered in the Password field. Be sure the password conforms to the allowed number of characters. The password characters are not displayed on the page, but are disguised in a browser-specific manner.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.10.2 Enable Password

Use the Enable Password Configuration page to configure the enable password. To access this page, click **System > Passwords > Enable Password**.

Figure 4.125 System > Passwords > Enable Password

The following table describes the items in the previous figure.

Item	Description
Enable Password	Specify the password all users must enter after executing the enable command at the CLI prompt.
Confirm Enable Password	Type the password again to confirm that you have entered it correctly.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.10.3 Password Rules

Use the Password Rules page to configure settings that apply to all user passwords. To access this page, click **System > Passwords > Password Rules**.

Figure 4.126 System > Passwords > Password Rules

The following table describes the items in the previous figure.

Item	Description
Minimum Length	The minimum number of characters required for a valid password.
Aging (Days)	The number of days that a user password is valid from the time the password is set. Once a password expires, the user is required to enter a new password at the next login.
History	The number of previous passwords that are retained to prevent password reuse. This helps to ensure that a user does not attempt to reuse the same password too often.
Lockout Attempts	The number of local authentication attempts that are allowed to fail before the user account is automatically locked.
Strength Check	Enables or disables the password strength checking feature. Enabling this feature forces the user to configure passwords that comply with the various strong password configuration parameters that are defined on this page.
Minimum Number of Uppercase Letters	The minimum number of upper-case letters that a valid password must contain.

Item	Description
Minimum Number of Lowercase Letters	The minimum number of lower-case letters that a valid password must contain.
Minimum Number of Numeric Characters	The minimum number of numeric characters that a valid password must contain.
Minimum Number of Special Characters	The minimum number of special characters (such as the keyboard symbols @, \$, &) that a valid password must contain.
Maximum Number of Repeated Characters	The maximum number of characters of any type that are allowed to repeat in a valid password. Repetition is defined as the same character occurring in succession anywhere within the password, such as "11" or "%%%" or "EEEE".
Maximum Number of Consecutive Characters	The maximum number of characters belonging to a sequence that are allowed to occur in a valid password. Consecutive characters are defined as a sequential pattern of case-sensitive alphabetic or numeric characters, such as "2345" or "def" or "YZ".
Minimum Character Classes	This minimum number of character classes, defined as the various password strength categories listed above, that must be met in order for a password to be considered valid. It is permissible, therefore, to define strength checking criteria for each of the different types of conditions, but only require a valid password to meet some of them. The number of these character classes that must be met is specified by this value.
Exclude Keyword Name	<p>The list of keywords that a valid password must not contain. Excluded keyword checking is case-insensitive. Additionally, a password cannot contain the backwards version of an excluded keyword. For example, if pass is an excluded keyword, passwords such as 23passA2c, ssapword, and PAsSworD are prohibited. Use the plus and minus buttons to perform the following tasks:</p> <ul style="list-style-type: none"> ■ To add a keyword to the list, click + button, type the word to exclude in the Exclude Keyword Name field, and click Submit. ■ To remove a keyword from the list, click - button associated with the keyword to remove and confirm the action. ■ To remove all keywords from the list, click - button in the header row and confirm the action.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.10.4 Last Password

Use the Last Password Result page view information about the last attempt to set a user password. If the password set was unsuccessful, a reason for the failure is given.

To access this page, click **System > Passwords > Last Password**.




Figure 4.127 System > Passwords > Last Password

The following table describes the items in the previous figure.

Item	Description
Last Result	Displays information about the last (User/Line/Enable) password configuration result. If the field is blank, no passwords have been configured on the device. Otherwise, the field shows that the password was successfully set or provides information about the type of password configuration that failed and why it could not be set.
Strength Check	Displays Enabled if Strength Check is applied in last password change, otherwise it displays Disabled.
Refresh	Click Refresh to update the screen.

4.3.10.5 Reset Passwords

Use the Reset Passwords page to reset all system login passwords to factory default values. When you click **Reset**, the reset action is initiated.

To access this page, click **System > Passwords > Reset Passwords**.



Figure 4.128 System > Passwords > Reset Passwords

The following table describes the items in the previous figure.

Item	Description
Reset	Click Reset to initiates a reset of all login passwords to their factory default setting after displaying a confirmation message.

4.3.11 Port

The pages in the Port folder allow you to view and monitor the physical port information for the ports available on the switch. The Port folder has links to the following pages:

4.3.11.1 Summary

Use the Port Summary page to view the settings for all physical ports on the platform.

To access this page, click **System > Port > Summary**.

Interface	Interface Index	Type	Admin Mode	Physical Mode	Physical Status	STP Mode	LACP Mode	Link Status
ge0/1	1	Normal	Enabled	Auto	100 Mbps Full Duplex	Enabled	Enabled	Link Up
ge0/2	2	Normal	Enabled	Auto		Enabled	Enabled	Link Down
ge0/3	3	Normal	Enabled	Auto		Enabled	Enabled	Link Down
ge0/4	4	Normal	Enabled	Auto		Enabled	Enabled	Link Down
ge0/5	5	Normal	Enabled	Auto		Enabled	Enabled	Link Down
ge0/6	6	Normal	Enabled	Auto		Enabled	Enabled	Link Down
ge0/7	7	Normal	Enabled	Auto		Enabled	Enabled	Link Down
ge0/8	8	Normal	Enabled	Auto		Enabled	Enabled	Link Down
ge0/9	9	Normal	Enabled	Auto		Enabled	Enabled	Link Down
ge0/10	10	Normal	Enabled	Auto		Enabled	Enabled	Link Down

Figure 4.129 System > Port > Summary

The following table describes the items in the previous figure.

Item	Description
Interface	Identifies the port or LAG.

Item	Description
Interface Index	The interface index object value assigned by the IF-MIB. This value is used to identify the interface when managing the device by using SNMP.
Type	The interface type, which is one of the following: <ul style="list-style-type: none"> ■ Normal: The port is a normal port, which means it is not a LAG member or configured for port mirroring. ■ Trunk Member: The port is a member of a LAG. ■ Mirrored: The port is configured to mirror its traffic (ingress, egress, or both) to another port (the probe port). ■ Probe: The port is configured to receive mirrored traffic from one or more source ports.
Admin Mode	The administrative mode of the interface. If a port or LAG is administratively disabled, it cannot forward traffic.
Physical Mode	The port speed and duplex mode. If the mode is Auto, the port's maximum capability are advertised, and the duplex mode and speed are set from the auto-negotiation process. The physical mode for a LAG is reported as "LAG".
Physical Status	Indicates the port speed and duplex mode for physical interfaces. The physical status for LAGs is not reported. When a port is down, the physical status is unknown.
STP Mode	The Spanning Tree Protocol (STP) Administrative Mode associated with the port or LAG. STP is a layer 2 protocol that provides a tree topology for switches on a bridged LAN. STP allows a network to have redundant paths without the risk of network loops. by providing a single path between end stations on a network. The possible values for STP mode are: <ul style="list-style-type: none"> ■ Enable: Spanning tree is enabled for this port. ■ Disable: Spanning tree is disabled for this port.
LACP Mode	Shows the administrative mode of the Link Aggregation Control Protocol (LACP), which is one of the following: <ul style="list-style-type: none"> ■ Enabled: The port uses LACP for dynamic LAG configuration. When LACP is enabled, the port sends and receives LACP PDUs with its link partner to confirm that the external switch is also configured for link aggregation. ■ Disabled: The port is supports static LAG configuration only. This mode might be used when the port is connected to a device that does not support LACP. When a port is added to a LAG as a static member, it neither transmits nor receives LACP PDUs.
Link Status	Indicates whether the link is up or down. The link is the physical connection between the port or LAG and the interface on another device.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.

4.3.11.2 Description

Use the Port Description page to configure a human-readable description of the port. To access this page, click **System > Port > Description**.

Interface	Physical Address	PortList Bit Offset	Interface Index	Port Description
ge0/1	00:11:22:33:44:55	1	1	
ge0/2	00:11:22:33:44:55	2	2	
ge0/3	00:11:22:33:44:55	3	3	
ge0/4	00:11:22:33:44:55	4	4	
ge0/5	00:11:22:33:44:55	5	5	
ge0/6	00:11:22:33:44:55	6	6	
ge0/7	00:11:22:33:44:55	7	7	
ge0/8	00:11:22:33:44:55	8	8	
ge0/9	00:11:22:33:44:55	9	9	
ge0/10	00:11:22:33:44:55	10	10	

Figure 4.130 System > Port > Description

The following table describes the items in the previous figure.

Item	Description
Interface	Identifies the port or LAG.
Physical Address	The MAC address of the interface.
PortList Bit Offset	The bit offset value that corresponds to the interface when the MIB object type Port List is used when managing the device by using SNMP.
Interface Index	The interface index object value assigned by the IF-MIB. This value is used to identify the interface when managing the device by using SNMP.
Port Description	The current description, if any, associated with the interface to help identify it.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.

4.3.11.3 Cable Test

The cable test feature enables you to determine the cable connection status on a selected port. You can also obtain an estimate of the length of the cable connected to the port, if the PHY on the ports supports this functionality.

Note! *The cable test feature is supported only for copper cable. It is not supported for optical fiber cable.*



To access this page, click **System > Port > Cable Test**.

Figure 4.131 System > Port > Cable Test

The following table describes the items in the previous figure.

Item	Description
Interface	Click the drop-down menu to select the port with the connected cable to test.

Item	Description
Failure Location Distance	The estimated distance from the end of the cable to the failure location. <i>NOTE: This field displays a value only when the Cable Status is Open or Short; otherwise, this field is blank.</i>
Cable Length (Meters)	The estimated length of the cable. If the cable length cannot be determined, Unknown is displayed. This field shows the range between the shortest estimated length and the longest estimated length. <i>NOTE: This field displays a value only when the Cable Status is Normal; otherwise, this field is blank.</i>
Cable Status	Displays the cable status as one of the following: <ul style="list-style-type: none"> ■ Normal: The cable is working correctly. ■ Open: The cable is disconnected, or there is a faulty connector. ■ Open and Short: There is an electrical short in the cable. ■ Cable status test failed: The cable status could not be determined. The cable may in fact be working.
Test Cable	Click Test Cable to perform a cable test on the selected interface. The cable test may take up to 2 seconds to complete. If the port has an active link, the link is not taken down, and the Cable Status always indicates Normal. The test returns a cable length estimate if this feature is supported by the PHY for the current link speed. <i>NOTE: If the link is down and a cable is attached to a 10/100 Ethernet adapter, the Cable Status may indicate Open or Short because some Ethernet adapters leave unused wire pairs unterminated or grounded.</i>

4.3.11.4 Mirroring

Port mirroring selects the network traffic for analysis by a network analyzer. This is done for specific ports of the switch. As such, many switch ports are configured as source ports and one switch port is configured as a destination port. You have the ability to configure how traffic is mirrored on a source port. Packets that are received on the source port, that are transmitted on a port, or are both received and transmitted, can be mirrored to the destination port.

The packet that is copied to the destination port is in the same format as the original packet on the wire. This means that if the mirror is copying a received packet, the copied packet is VLAN tagged or untagged as it was received on the source port. If the mirror is copying a transmitted packet, the copied packet is VLAN tagged or untagged as it is being transmitted on the source port.

Use the Multiple Port Mirroring page to define port mirroring sessions.

To access this page, click **System > Port > Mirroring**.





Figure 4.132 System > Port > Mirroring

The following table describes the items in the previous figure.

Item	Description
Session ID	The port mirroring session ID. The number of sessions allowed is platform specific.
Mode	The administrative mode for the selected port mirroring session. If the mode is disabled, the configured source is not mirroring traffic to the destination.
Destination	<p>The interface that receives traffic from all configured source ports.</p> <p>After you click  button, the Destination Configuration window opens. The following information describes the additional fields available in this window.</p> <ul style="list-style-type: none"> ■ Type: The type of interface to use as the destination, which is one of the following: <ul style="list-style-type: none"> – None: The destination is not configured. – Remote VLAN: Traffic is mirrored to the VLAN on the system that is configured as the RSPAN VLAN. In an RSPAN configuration, the destination should be the Remote VLAN on any device that does not have a port connected to the network traffic analyzer. – Interface: Traffic is mirrored to a physical port on the local device. The interface is the probe port that is connected to a network traffic analyzer. ■ Remote VLAN: The VLAN that is configured as the RSPAN VLAN. ■ Port: The port to which traffic is mirrored. If the Type is Remote VLAN, the selected port is a reflector port. The reflector port is a trunk port that carries the mirrored traffic towards the destination device. If the Type is Interface, the selected port is the probe port that is connected to a network traffic analyzer.
Source	The ports or VLAN configured to mirror traffic to the destination. You can configure multiple source ports or one source VLAN per session. The source VLAN can also be a remote VLAN.
Direction	<p>The direction of traffic on the source port (or source ports) or VLAN that is sent to the specified destination. A source VLAN mirrors all received and transmitted packets to the destination. Possible values for source ports are:</p> <ul style="list-style-type: none"> ■ Tx and Rx: Both ingress and egress traffic. ■ Rx: Ingress traffic only. ■ Tx: Egress traffic only.
Refresh	Click Refresh to update the screen.
Configure Session	Click Configure Session to configure the administrative mode for a port mirroring session or to select an ACL for flow-based mirroring.
Configure Source	Click Configure Source to configure one or more source ports or a VLAN for the mirroring session and to determine which traffic is mirrored (Tx, Rx, or both).
Remove Source	Click Remove Source to remove the selected source ports.

4.3.11.5 Mirroring Summary

To access this page, click **System > Port > Mirroring Summary**.

Session ID	Admin Mode	Probe Port	Src VLAN	Mirrored Port	Reflector Port	Src RVLAN	Dst RVLAN	Type
1	Disabled							
2	Disabled							
3	Disabled							
4	Disabled							

Figure 4.133 System > Port > Mirroring Summary

The following table describes the items in the previous figure.

Item	Description
Session ID	The port mirroring session ID. The number of sessions allowed is platform specific.
Admin Mode	The administrative mode for the selected port mirroring session. If the mode is disabled, the configured source is not mirroring traffic to the destination.
Probe Port	The interface that receives traffic from all configured source ports.
Src VLAN	The VLAN configured to mirror traffic to the destination. You can configure one source VLAN per session. The source VLAN can also be a remote VLAN.
Mirrored Port	The ports configured to mirror traffic to the destination. You can configure multiple source ports per session.
Reflector Port	This port carries all the mirrored traffic at source switch.
Src RVLAN	The VLAN configured as the RSPAN VLAN is the source. In an RSPAN configuration, the remote VLAN is the source on the destination device that has a physical port connected to the network traffic analyzer.
Dst RVLAN	Traffic is mirrored to the VLAN on the system that is configured as the RSPAN VLAN. In an RSPAN configuration, the destination should be the Remote VLAN on any device that does not have a port connected to the network traffic analyzer.
Type	The type of traffic on the source port (or source ports) or VLAN that is sent to the specified destination. A source VLAN mirrors all received and transmitted packets to the destination. Possible values for source ports are: <ul style="list-style-type: none"> ■ Tx and Rx: Both ingress and egress traffic. ■ Rx: Ingress traffic only. ■ Tx: Egress traffic only.
Refresh	Click Refresh to update the screen.

4.3.12 Slot

4.3.12.1 Configuration

Use the Configuration page to view information about the cards installed in the device's slots and to configure settings for the slots available on the device. Support for adding cards to a slot or changing the slot configuration is platform dependent.

To access this page, click **System > Slot > Configuration**.

Slot	Status	Administrative State	Power State	Card Model	Card Description
0	Full	Enable	Enable	BCM56150	Broadcom BCM56150 - 24 GE + 4 XE Ethernet Line Card

Figure 4.134 System > Slot > Configuration

The following table describes the items in the previous figure.

Item	Description
Slot	Identifies the slot number.
Status	Indicates whether the slot is empty or full.
Administrative State	Indicates whether the slot is administratively enabled or disabled. For some devices, you can change the Administrative State when you add or edit slot information.
Power State	Indicates whether the device is providing power to the slot. For some devices, you can change the Power State when you add or edit slot information.
Card Model	The model ID of the card configured for the slot.
Card Description	The description of the card configured for the slot.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new card. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

To add a new card:

Click **System > Slot > Configuration > Add**.

Add New Card

Unit	1
Slot	0
Card Index	1
Card Description	BCM56150
Status	Full
Administrative State	<input checked="" type="radio"/> Enable <input type="radio"/> Disable
Power State	<input checked="" type="radio"/> Enable <input type="radio"/> Disable
Inserted Card Model	BCM56150
Inserted Card Description	Broadcom BCM56150 - 24 GE + 4 XE Ethernet Line Card
Configured Card Model	BCM56150
Configured Card Description	Broadcom BCM56150 - 24 GE + 4 XE Ethernet Line Card
Pluggable	False
Power Down	False

Submit Cancel

Figure 4.135 System > Slot > Configuration > Add

The following table describes the items in the previous figure.

Item	Description
Unit	Identifies the unit number of the device (in the stack of devices) on which to add the new card.

Item	Description
Card Index	Identifies the index number assigned to the card. This value is helpful when configuring the system by using SNMP.
Inserted Card Model	The model ID of the card plugged into the slot.
Inserted Card Description	The description of the card plugged into the slot.
Configured Card Model	The model ID of the card configured for the slot.
Configured Card Description	The description of the card configured for the slot.
Pluggable	If the value is True, the card can be administratively enabled or disabled. If the value is False, the Administrative State cannot be configured.
Power Down	If the value is True, the Power State can be administratively enabled or disabled. If the value is False, the Power State cannot be configured.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.3.12.2 Supported Cards

To access this page, click **System > Slot > Supported Cards**.

Card Index	Supported Cards	Card Type	Card Model	Card Description
1	BCM5G150	0x56150101	BCM5G150	Broadcom BCM56150 - 24 GE + 4 XE Ethernet Line Card
2	BCM53344	0x53344101	BCM53344	Broadcom BCM53344 - 24 GE + 4 GE Ethernet Line Card
3	BCM53344	0x53344102	BCM53344	Broadcom BCM53344 - 24 GE + 4 GE Ethernet PoE Line Card
4	BCM53343	0x53343101	BCM53343	Broadcom BCM53343 - 16 GE Line Card

Figure 4.136 System > Slot > Supported Cards

The following table describes the items in the previous figure.

Item	Description
Card Index	The index assigned to the card type.
Supported Cards	The model of the card that can be supported.
Card Type	The hardware type of the supported card, which is assigned by the manufacturer.
Card Model	Similar to the Supported Cards information, this field identifies the model of the supported card.
Card Description	Description of the supported card, which might include the manufacturer's product number and information about number and speed of the supported interfaces.
Refresh	Click Refresh to update the screen.

4.3.13 Statistics

4.3.13.1 System

Switch

The Switch Statistics page shows summary information about traffic transmitted and received on the device, entries in the MAC address table, and Virtual Local Area Networks (VLANs) that exist on the device.

To access this page, click **System > Statistics > System > Switch**.

Statistics	Transmit	Receive
Octets Without Error	5675064	3540588
Packets Without Errors	9643	17980
Packets Discarded	0	0
Unicast Packets	9529	8423
Multicast Packets	111	2675
Broadcast Packets	3	6882

Status	FDB Entries	VLANs
Current Usage	10	1
Peak Usage	16	1
Maximum Allowed	16384	4093
Static Entries	1	1
Dynamic Entries	9	0
Total Entries Deleted	N/A	0

System	
Interface	29
Time Since Counters Last Cleared	0d02:55:17

Figure 4.137 System > Statistics > System > Switch

The following table describes the items in the previous figure.

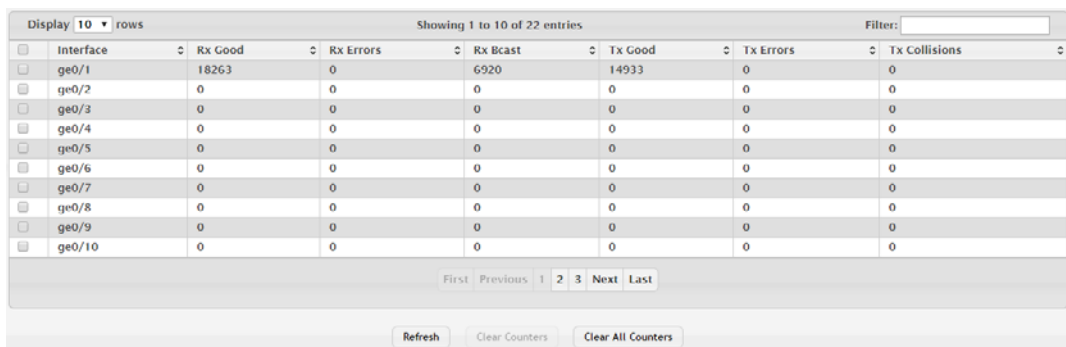
Item	Description
Statistics	
Octets Without Error	The total number of octets (bytes) of data successfully transmitted or received by the processor (excluding framing bits but including FCS octets).
Packets Without Errors	The total number of packets including unicast, broadcast, and multicast packets, successfully transmitted or received by the processor.
Packets Discarded	The number of outbound (Transmit column) or inbound (Receive column) packets which were chosen to be discarded even though no errors had been detected to prevent their being deliverable to a higher-layer protocol. A possible reason for discarding a packet could be to free up buffer space.
Unicast Packets	The number of subnetwork-unicast packets delivered to or received from a higher-layer protocol.
Multicast Packets	The total number of packets transmitted or received by the device that were directed to a multicast address. Note that this number does not include packets directed to the broadcast address.
Broadcast Packets	The total number of packets transmitted or received by the device that were directed to the broadcast address. Note that this number does not include multicast packets.
Status	
Current Usage	In the FDB Entries column, the value shows the number of learned and static entries in the MAC address table. In the VLANs column, the value shows the total number of static and dynamic VLANs that currently exist in the VLAN database.
Peak Usage	The highest number of entries that have existed in the MAC address table or VLAN database since the most recent reboot.
Maximum Allowed	The maximum number of statically configured or dynamically learned entries allowed in the MAC address table or VLAN database.

Item	Description
Static Entries	The current number of entries in the MAC address table or VLAN database that an administrator has statically configured.
Dynamic Entries	The current number of entries in the MAC address table or VLAN database that have been dynamically learned by the device.
Total Entries Deleted	The number of VLANs that have been created and then deleted since the last reboot. This field does not apply to the MAC address table entries.
System	
Interface	The interface index object value of the interface table entry associated with the Processor of this switch. This value is used to identify the interface when managing the device by using SNMP.
Time Since Counters Last Cleared	The amount of time in days, hours, minutes, and seconds, that has passed since the statistics for this device were last reset.
Refresh	Click Refresh to update the screen.
Clear Counters	Click Clear Counters to reset all counters to zero. <i>NOTE: The Packets Discarded cannot be cleared.</i>

Port Summary

The Port Summary Statistics page shows statistical information about the packets received and transmitted by each port and LAG.

To access this page, click **System > Statistics > System > Port Summary**.



Interface	Rx Good	Rx Errors	Rx Bcast	Tx Good	Tx Errors	Tx Collisions
ge0/1	18263	0	6920	14933	0	0
ge0/2	0	0	0	0	0	0
ge0/3	0	0	0	0	0	0
ge0/4	0	0	0	0	0	0
ge0/5	0	0	0	0	0	0
ge0/6	0	0	0	0	0	0
ge0/7	0	0	0	0	0	0
ge0/8	0	0	0	0	0	0
ge0/9	0	0	0	0	0	0
ge0/10	0	0	0	0	0	0

Figure 4.138 System > Statistics > System > Port Summary

The following table describes the items in the previous figure.

Item	Description
Interface	Identifies the port or LAG.
Rx Good	The total number of inbound packets received by the interface without errors.
Rx Errors	The number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol.
Rx Bcast	The total number of good packets received that were directed to the broadcast address. Note that this number does not include multicast packets.
Tx Good	The total number of outbound packets transmitted by the interface to its Ethernet segment without errors.
Tx Errors	The number of outbound packets that could not be transmitted because of errors.
Tx Collisions	The best estimate of the total number of collisions on this Ethernet segment.
Refresh	Click Refresh to update the screen.

Item	Description
Clear Counters	Click Clear Counters to reset the selected counters to zero.
Clear All Counters	Click Clear All Counters to reset all counters to zero.

Port Detailed

The Port Detailed Statistics page shows detailed information about the traffic transmitted and received by each interface.

To access this page, click **System > Statistics > System > Port Detailed**.

Interface	0/1 ▾	
Maximum Frame Size	1518	
MTU	1500	
Packet Lengths Received and Transmitted		
64 Octets	0	
65-127 Octets	0	
128-255 Octets	0	
256-511 Octets	0	
512-1023 Octets	0	
1024-1518 Octets	0	
1519-1522 Octets	0	
1523-2047 Octets	0	
2048-4095 Octets	0	
4096-9216 Octets	0	
Basic		
	Transmit	Receive
Unicast Packets	0	0
Multicast Packets	0	0
Broadcast Packets	0	0
Total Packets (Octets)	0	0
Packets > 1518 Octets	0	0
802.3x Pause Frames	0	0
FCS Errors	0	0
Protocol		
	Transmit	Receive
STP BPDU	0	0
RSTP BPDU	0	0
MSTP BPDU	0	0
SSTP BPDU	0	0
GVRP PDU	0	0
GMRP PDU	0	0
EAPOL Frames	0	0
Advanced - Transmit		
Total Transmit Packets Discarded	0	
Single Collision Frames	0	
Multiple Collision Frames	0	

Figure 4.139 System > Statistics > System > Port Detailed

The following table describes the items in the previous figure.

Item	Description
Interface	Identifies the port or LAG. To view the statistics for a specific interface, select the interface number from the drop-down menu. The page automatically refreshes with the statistics for the selected interface.
Maximum Frame Size	The maximum Ethernet frame size the interface supports or is configured to support. The maximum frame size includes the Ethernet header, CRC, and payload.
MTU	Indicates MTU (Maximum Transmit Unit) of the interface. The actual frame size is calculated by adding ethernet header size in MTU.
Packet Lengths Received and Transmitted	
64 Octets	The total number of packets (including bad packets) received or transmitted that were 64 octets in length (excluding framing bits but including FCS octets).
65-127 Octets	The total number of packets (including bad packets) received or transmitted that were between 65 and 127 octets in length inclusive (excluding framing bits but including FCS octets).
128-255 Octets	The total number of packets (including bad packets) received or transmitted that were between 128 and 255 octets in length inclusive (excluding framing bits but including FCS octets).
256-511 Octets	The total number of packets (including bad packets) received or transmitted that were between 256 and 511 octets in length inclusive (excluding framing bits but including FCS octets).

Item	Description
512-1023 Octets	The total number of packets (including bad packets) received or transmitted that were between 512 and 1023 octets in length inclusive (excluding framing bits but including FCS octets).
1024-1518 Octets	The total number of packets (including bad packets) received or transmitted that were between 1024 and 1518 octets in length inclusive (excluding framing bits but including FCS octets).
1519-1522 Octets	The total number of packets (including bad packets) received or transmitted that were between 1519 and 1522 octets in length inclusive (excluding framing bits but including FCS octets).
1523-2047 Octets	The total number of packets (including bad packets) received or transmitted that were between 1523 and 2047 octets in length inclusive (excluding framing bits but including FCS octets).
2048-4095 Octets	The total number of packets (including bad packets) received or transmitted that were between 2048 and 4095 octets in length inclusive (excluding framing bits but including FCS octets).
4096-9216 Octets	The total number of packets (including bad packets) received or transmitted that were between 4096 and 9216 octets in length inclusive (excluding framing bits but including FCS octets).
Basic	
Unicast Packets	The Transmit column shows the total number of packets that higher-level protocols requested be transmitted to a subnetwork unicast address, including those that were discarded or not sent. The Receive column shows the number of subnetwork unicast packets delivered to a higher-layer protocol.
Multicast Packets	The Transmit column shows the total number of packets that higher-level protocols requested be transmitted to a multicast address, including those that were discarded or not sent. The Receive column shows the number of multicast packets delivered to a higher-layer protocol.
Broadcast Packets	The Transmit column shows the total number of packets that higher-level protocols requested be transmitted to a broadcast address, including those that were discarded or not sent. The Receive column shows the number of broadcast packets delivered to a higher-layer protocol.
Total Packets (Octets)	The total number of octets of data (including those in bad packets) transmitted or received on the interface (excluding framing bits but including FCS octets). This object can be used as a reasonable estimate of Ethernet utilization. If greater precision is desired, the etherStatsPkts and etherStatsOctets objects should be sampled before and after a common interval.
Packets > 1518 Octets	The total number of packets transmitted or received by this interface that were longer than 1518 octets (excluding framing bits, but including FCS octets) and were otherwise well formed. This counter has a maximum increment rate of 815 counts per sec at 10 Mb/s.
802.3x Pause Frames	The number of MAC Control frames transmitted or received by this interface with an opcode indicating the PAUSE operation. This counter does not increment when the interface is operating in half-duplex mode.
FCS Errors	The total number of packets transmitted or received by this interface that had a length (excluding framing bits, but including FCS octets) between 64 and 1518 octets, inclusive, but had a bad Frame Check Sequence (FCS) with an integral number of octets.

Item	Description
Protocol	
STP BPDUs	The number of Spanning Tree Protocol (STP) Bridge Protocol Data Units (BPDUs) transmitted or received by the interface.
RSTP BPDUs	The number of Rapid STP BPDUs transmitted or received by the interface.
MSTP BPDUs	The number of Multiple STP BPDUs transmitted or received by the interface.
SSTP BPDUs	The number of Shared Spanning Tree Protocol (SSTP) Bridge Protocol Data Units (BPDUs) transmitted or received by the interface.
GVRP PDUs	The number of Generic Attribute Registration Protocol (GARP) VLAN Registration Protocol (GVRP) PDUs transmitted or received by the interface.
GMRP PDUs	The number of GARP Multicast Registration Protocol (GMRP) PDUs transmitted or received by the interface.
EAPOL Frames	The number of Extensible Authentication Protocol (EAP) over LAN (EAPOL) frames transmitted or received by the interface for IEEE 802.1X port-based network access control.
Advanced - Transmit	
Total Transmit Packets Discarded	The sum of single collision frames discarded, multiple collision frames discarded, and excessive frames discarded.
Single Collision Frames	A count of the number of successfully transmitted frames on a particular interface for which transmission is inhibited by exactly one collision.
Multiple Collision Frames	A count of the number of successfully transmitted frames on a particular interface for which transmission is inhibited by more than one collision.
Excessive Collision Frames	A count of frames for which transmission on a particular interface fails due to excessive collisions.
Underrun Errors	The total number of frames discarded because the transmit FIFO buffer became empty during frame transmission.
GMRP Failed Registrations	The number of times attempted GMRP registrations could not be completed.
GVRP Failed Registrations	The number of times attempted GVRP registrations could not be completed.
Percent Utilization Transmitted	The value of link utilization in percentage representation for TX line.
Advanced - Receive	
Total Packets Received Not Forwarded	The number of inbound packets which were chosen to be discarded to prevent them from being delivered to a higher-layer protocol, even though no errors had been detected. One possible reason for discarding such a packet is to free up buffer space.
Total Packets Received With MAC Errors	The total number of inbound packets that contained errors preventing them from being delivered to a higher-layer protocol.
Overruns	The total number of frames discarded as this port was overloaded with incoming packets, and could not keep up with the inflow.
Alignment Errors	The total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, but had a bad Frame Check Sequence (FCS) with a non-integral number of octets.

Item	Description
Jabbers Received	The total number of packets received that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error). Note that this definition of jabber is different than the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition where any packet exceeds 20 ms. The allowed range to detect jabber is between 20 ms and 150 ms.
Fragments Received	The total number of packets received that were less than 64 octets in length with ERROR CRC (excluding framing bits but including FCS octets).
Undersize Received	The total number of packets received that were less than 64 octets in length with GOOD CRC (excluding framing bits but including FCS octets).
Unacceptable Frame Type	The number of frames discarded from this interface due to being a frame type that the interface cannot accept.
Percent Utilization Received	The value of link utilization in percentage representation for RX line.
Time Since Counters Last Cleared	The amount of time in days, hours, minutes, and seconds, that has passed since the statistics for this interface were last reset.
Refresh	Click Refresh to update the screen.
Clear Counters	Click Clear Counters to reset the detailed statistics for the selected interface to the default values.
Clear All Counters	Click Clear All Counters to reset the detailed statistics for all interfaces to the default values.

Network DHCPv6

The Network Port DHCPv6 Client Statistics page displays the DHCPv6 client statistics values for the network interface. The DHCPv6 client on the device exchanges several different types of UDP messages with one or more network DHCPv6 servers during the process of acquiring address, prefix, or other relevant network configuration information from the server. The values indicate the various counts that have accumulated since they were last cleared.

To access this page, click **System > Statistics > System > Network DHCPv6**.

Advertisement Packets Received	0
Reply Packets Received	0
Received Advertisement Packets Discarded	0
Received Reply Packets Discarded	0
Malformed Packets Received	0
Total Packets Received	0
Solicit Packets Transmitted	100
Request Packets Transmitted	0
Renew Packets Transmitted	0
Rebind Packets Transmitted	0
Release Packets Transmitted	0
Total Packets Transmitted	100

Refresh Clear Counters

Figure 4.140 System > Statistics > System > Network DHCPv6

The following table describes the items in the previous figure.

Item	Description
Advertisement Packets Received	Number of DHCPv6 advertisement messages received from one or more DHCPv6 servers in response to the client's solicit message.
Reply Packets Received	Number of DHCPv6 reply messages received from one or more DHCPv6 servers in response to the client's request message.

Item	Description
Received Advertisement Packets Discarded	Number of DHCPv6 advertisement messages received from one or more DHCPv6 servers to which the client did not respond.
Received Reply Packets Discarded	Number of DHCPv6 reply messages received from one or more DHCPv6 servers to which the client did not respond.
Malformed Packets Received	Number of messages received from one or more DHCPv6 servers that were improperly formatted.
Total Packets Received	Total number of messages received from all DHCPv6 servers.
Solicit Packets Transmitted	Number of DHCPv6 solicit messages the client sent to begin the process of acquiring network information from a DHCPv6 server.
Request Packets Transmitted	Number of DHCPv6 request messages the client sent in response to a DHCPv6 server's advertisement message.
Renew Packets Transmitted	Number of renew messages the DHCPv6 client has sent to the server to request an extension of the lifetime of the information provided by the server. This message is sent to the DHCPv6 server that originally assigned the addresses and configuration information.
Rebind Packets Transmitted	Number of rebind messages the DHCPv6 client has sent to any available DHCPv6 server to request an extension of its addresses and an update to any other relevant information. This message is sent only if the client does not receive a response to the renew message.
Release Packets Transmitted	Number of release messages the DHCPv6 client has sent to the server to indicate that it no longer needs one or more of the assigned addresses.
Total Packets Transmitted	Total number of messages sent to all DHCPv6 servers.
Refresh	Click Refresh to update the screen.
Clear Counters	Click Clear Counters to reset all counters to zero.

4.3.13.2 Time Based Group

Use the Time Based Group Statistics page to define criteria for collecting time-based statistics for interface traffic. The time-based statistics can be useful for troubleshooting and diagnostics purposes. The statistics application uses the system clock for time-based reporting, so it is important to configure the system clock (manually or through SNTP) before using this feature.

To access this page, click **System > Statistics > Time Based > Group**.




Figure 4.141 System > Statistics > Time Based > Group

The following table describes the items in the previous figure.

Item	Description
Group	<p>The type of traffic statistics to collect for the group, which is one of the following:</p> <ul style="list-style-type: none"> ■ Received: The number of packets received on the interfaces within the group. ■ Received Errors: The number of packets received with errors on the interfaces within the group. ■ Transmitted: The number of packets transmitted by the interfaces within the group. ■ Received Transmitted: The number of packets received and transmitted by the interfaces within the group. ■ Port Utilization: The percentage of total bandwidth used by the port within the specified time period. ■ Congestion: The percentage of time within the specified time range that the ports experienced congestion.
Time Range	<p>The name of the periodic or absolute time range to use for data collection. The time range is configured by using the Time Range Summary and Time Range Entry Summary pages. The time range must be configured on the system before the time-based statistics can be collected.</p>
Reporting Methods	<p>The methods for reporting the collected statistics at the end of every configured time range interval. The available options are:</p> <ul style="list-style-type: none"> ■ None: The statistics are not reported to the console or an external server. They can be viewed only by using the web interface or by issuing a CLI command. ■ Console: The statistics are displayed on the console. ■ E-Mail: The statistics are sent to an e-mail address. The SMTP server and e-mail address information is configured by using the appropriate Email Alerts pages. ■ Syslog: The statistics are sent to a remote syslog server. The syslog server information is configured on the Logging Hosts page.
Interfaces	<p>The interface or interfaces on which data is collected. To select multiple interfaces when adding a new group, CTRL + click each interface to include in the group.</p>
Refresh	<p>Click Refresh to update the screen.</p>
Add	<p>Click Add to add a new time based group. See the following procedure.</p>
Remove	<p>Click Remove to remove the selected entries.</p>

To add a new time based group:

Click **System > Statistics > Time Based > Group > Add**.

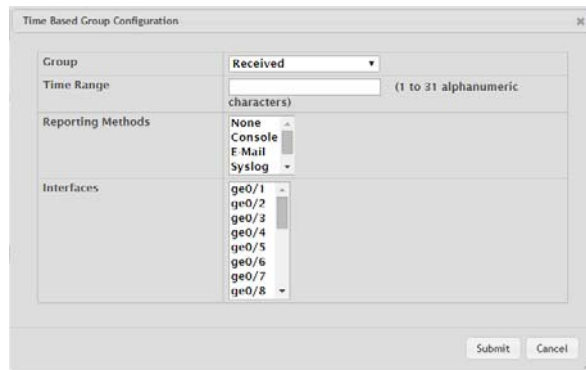


Figure 4.142 System > Statistics > Time Based > Group > Add

The following table describes the items in the previous figure.

Item	Description
Group	<p>The type of traffic statistics to collect for the group, which is one of the following:</p> <ul style="list-style-type: none"> ■ Received: The number of packets received on the interfaces within the group. ■ Received Errors: The number of packets received with errors on the interfaces within the group. ■ Transmitted: The number of packets transmitted by the interfaces within the group. ■ Received Transmitted: The number of packets received and transmitted by the interfaces within the group. ■ Port Utilization: The percentage of total bandwidth used by the port within the specified time period. ■ Congestion: The percentage of time within the specified time range that the ports experienced congestion.
Time Range	<p>The name of the periodic or absolute time range to use for data collection. The time range is configured by using the Time Range Summary and Time Range Entry Summary pages. The time range must be configured on the system before the time-based statistics can be collected.</p>
Reporting Methods	<p>The methods for reporting the collected statistics at the end of every configured time range interval. The available options are:</p> <ul style="list-style-type: none"> ■ None: The statistics are not reported to the console or an external server. They can be viewed only by using the web interface or by issuing a CLI command. ■ Console: The statistics are displayed on the console. ■ E-Mail: The statistics are sent to an e-mail address. The SMTP server and e-mail address information is configured by using the appropriate Email Alerts pages. ■ Syslog: The statistics are sent to a remote syslog server. The syslog server information is configured on the Logging Hosts page.
Interfaces	<p>The interface or interfaces on which data is collected. To select multiple interfaces when adding a new group, CTRL + click each interface to include in the group.</p>
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

Flow Based

Use the Time Based Flow Statistics page to define criteria for collecting time-based statistics for specific traffic flows. The statistics include a per-interface hit count based on traffic that meets the match criteria configured in a rule for the interfaces included in the rule. The hit count statistics are collected only during the specified time range. The statistics application uses the system clock for time-based reporting. Configure the system clock (manually or through SNTP) before using the time-based statistics feature.

To access this page, click **System > Statistics > Time Based > Flow Based**.




Figure 4.143 System > Statistics > Time Based > Flow Based

The following table describes the items in the previous figure.

Item	Description
Reporting Methods	The methods for reporting the collected statistics at the end of every configured interval. To change the reporting methods for all flow-based statistics rules, click the Edit icon and select one or more methods. To reset the field to the default value, click the Reset icon. The available reporting methods are: <ul style="list-style-type: none">■ None: The statistics are not reported to the console or an external server. They can be viewed only by using the web interface or by issuing a CLI command.■ Console: The statistics are displayed on the console.■ E-Mail: The statistics are sent to an e-mail address. The SNTP server and e-mail address information is configured by using the appropriate Email Alerts pages.■ Syslog: The statistics are sent to a remote syslog server. The syslog server information is configured on the Logging Hosts page.
Rule Id	The number that identifies the flow-based statistics collection rule.
Time Range	The name of the periodic or absolute time range to use for data collection. The time range is configured by using the Time Range Summary and Time Range Entry Summary pages. The time range must be configured on the system before the time-based statistics can be collected.
Match Conditions	The criteria a packet must meet to match the rule.
Interfaces	The interface or interfaces on which the flow-based rule is applied. Only traffic on the specified interfaces is checked against the rule.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new time based flow. See the following procedure.
Remove	Click Remove to remove the selected entries.

To add a new time based flow:

Click **System > Statistics > Time Based > Flow Based > Add**.

Figure 4.144 System > Statistics > Time Based > Flow Based > Add

The following table describes the items in the previous figure.

Item	Description
Rule Id	The number that identifies the flow-based statistics collection rule.
Time Range	The name of the periodic or absolute time range to use for data collection. The time range is configured by using the Time Range Summary and Time Range Entry Summary pages. The time range must be configured on the system before the time-based statistics can be collected.
Interface	The interface or interfaces on which the flow-based rule is applied. Only traffic on the specified interfaces is checked against the rule.
Match Criteria	
Match All	Select this option to indicate that all traffic matches the rule and is counted in the statistics. This option is exclusive to all other match criteria, so if Match All is selected, no other match criteria can be configured.
Source IP	The source IP address to match in the IPv4 packet header.
Destination IP	The destination IP address to match in the IPv4 packet header.
Source MAC	The source MAC address to match in the ingress frame header.
Destination MAC	The destination MAC address to match in the ingress frame header.
Source TCP Port	The TCP source port to match in the TCP header.
Destination TCP Port	The TCP destination port to match in the TCP header.
Source UDP Port	The UDP source port to match in the UDP header.
Destination UDP Port	The UDP destination port to match in the UDP header.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

Statistics

Use the Time Based Statistics page to view time-based statistics collected for the configured traffic groups and flow-based rules.

To access this page, click **System > Statistics > Time Based > Statistics**.




Figure 4.145 System > Statistics > Time Based > Statistics

The following table describes the items in the previous figure.

Item	Description
ID	The traffic group name or flow-based rule ID associated with the rest of the statistics in the row.
Interface	The interface on which the statistics were reported.
Counter Id	For traffic group statistics, this field identifies the type of traffic.
Counter Value	For traffic group statistics, this field shows the number of packets of the type identified by the Counter Id field that were reported on the interface during the time range.
Port Utilization	For a port utilization traffic group, this field reports the percentage of the total available bandwidth used on the interface during the time range.
Hit Count	For flow-based statistics, this field reports the number of packets that matched the flow-based rule criteria during the time range.
Refresh	Click Refresh to update the screen.

4.3.14 Status

4.3.14.1 ARP Cache

The ARP cache is a table maintained locally in each station on a network. ARP cache entries are learned by examining the source information in the ARP packet payload fields, regardless of whether it is an ARP request or response. Thus, when an ARP request is broadcast to all stations on a LAN segment or virtual LAN (VLAN), every recipient has the opportunity to store the sender's IP and MAC address in their respective ARP cache. The ARP response, being unicast, is normally seen only by the requestor, who stores the sender information in its ARP cache. Newer information always replaces existing content in the ARP cache.

The ARP cache can support 1024 entries, although this size is user-configurable to any value less than 1024. When multiple network interfaces are supported by a device, as is typical of a router, either a single ARP cache is used for all interfaces, or a separate cache is maintained per interface. While the latter approach is useful when network addressing is not unique per interface, this is not the case for Ethernet MAC address assignment so a single ARP cache is employed.

To access this page, click **System > Status > ARP Cache**.




Figure 4.146 System > Status > ARP Cache

The following table describes the items in the previous figure.

Item	Description
MAC Address	The physical (MAC) address associated with the IP address of the connection.
IP Address	The Internet (IP) address of the connection.
Interface	Shows the switch port through which the connection was established, or displays as Management if the connection occurred via a non-network port interface (if applicable).
Refresh	Click Refresh to update the screen.
Clear Entries	Click Clear Entries to clear all entries from the system ARP Cache.

4.3.14.2 Resource Status

Use the System Resource Status page to display the following memory information for the switch:

- Free memory
- Allocated memory
- CPU utilization by task
- Total CPU utilization at the following intervals:
 - Five seconds
 - One minute
 - Five minutes

To access this page, click **System > Status > Resource Status**.

The screenshot displays the 'System > Status > Resource Status' page. It features two main sections: 'Memory Usage' and 'CPU Utilization Report'. The 'Memory Usage' section shows 'Free Memory (Kbytes)' at 309360 and 'Alloc Memory (Kbytes)' at 189688. The 'CPU Utilization Report' section shows a table of tasks with columns for 'Task ID', 'Task Name', and CPU utilization percentages for '5 Seconds', '60 Seconds', and '300 Seconds'. The table lists tasks such as (ksoftirqd/0), (procmgr), (syncdb), proclOG, osapiTimer, bcmINTR, socdmadesc.0, bcmMEM_SCAN.0, bcmL2X.0, and bcmCNTR.0. A 'Refresh' button is located at the bottom of the table.

Memory Usage	
Free Memory (Kbytes)	309360
Alloc Memory (Kbytes)	189688

CPU Utilization Report					
Display 10 rows		Showing 1 to 10 of 38 entries			Filter:
Task ID	Task Name	5 Seconds	60 Seconds	300 Seconds	
3	(ksoftirqd/0)	0.00%	0.10%	0.09%	
1256	(procmgr)	1.01%	0.96%	1.01%	
1352	(syncdb)	0.00%	0.06%	0.11%	
1386	proclOG	0.40%	0.28%	0.27%	
1388	osapiTimer	1.41%	1.69%	1.70%	
1439	bcmINTR	0.00%	0.11%	0.09%	
1440	socdmadesc.0	0.00%	0.06%	0.13%	
1441	bcmMEM_SCAN.0	0.00%	0.23%	0.22%	
1443	bcmL2X.0	6.06%	6.26%	6.46%	
1444	bcmCNTR.0	0.60%	0.68%	0.70%	

Figure 4.147 System > Status > Resource Status

The following table describes the items in the previous figure.

Item	Description
Memory Usage	
Free Memory (Kbytes)	The amount of system memory that is currently available for allocation, specified in kilobytes.
Alloc Memory (Kbytes)	The amount of system memory that is currently allocated for use, specified in kilobytes.
CPU Utilization Report	
Task ID	System task identifier. The entry named Total represents the total CPU utilization, expressed as a percentage, that is used by the entire system for each of the specified time intervals.
Task Name	System task name.
5 Seconds	The percentage amount of CPU utilization consumed by the corresponding task in the last 5 seconds.

Item	Description
60 Seconds	The percentage amount of CPU utilization consumed by the corresponding task in the last 60 seconds.
300 Seconds	The percentage amount of CPU utilization consumed by the corresponding task in the last 300 seconds.
Refresh	Click Refresh to update the screen.

4.3.14.3 Resource Configuration

Use the System Resource Configuration page to configure the threshold parameters for monitoring CPU utilization and the amount of free memory in the system.

To access this page, click **System > Status > Resource Configuration**.

Note! *Setting any these configuration values to zero disables monitoring of that particular item and suppresses its corresponding event notification.*



Rising Threshold (%)	0	(0 to 100, 0 = Default, 0 = Disable)
Rising Threshold Interval (Seconds)	0	(0 to 86400, 0 = Default, 0 = Disable) Multiple of 5
Falling Threshold (%)	0	(0 to 100, 0 = Default, 0 = Disable)
Falling Threshold Interval (Seconds)	0	(0 to 86400, 0 = Default, 0 = Disable) Multiple of 5
Free Memory Threshold (Kbytes)	0	(0 to 499048, 0 = Default, 0 = Disable)

Figure 4.148 System > Status > Resource Configuration

The following table describes the items in the previous figure.

Item	Description
Rising Threshold (%)	The CPU utilization rising threshold, expressed as a percentage. When the CPU utilization is increasing, an event is signaled when it reaches or exceeds this level.
Rising Threshold Interval (Seconds)	The CPU utilization rising threshold interval in seconds. This represents how often the current CPU utilization is checked against the configured rising threshold value.
Falling Threshold (%)	The CPU utilization falling threshold, expressed as a percentage. When the CPU utilization is decreasing, an event is signaled when it reaches or falls below this level.
Falling Threshold Interval (Seconds)	The CPU utilization falling threshold interval in seconds. This represents how often the current CPU utilization is checked against the configured falling threshold value.
Free Memory Threshold (Kbytes)	The free memory threshold in kilobytes. If enabled, an event is signaled when the amount of free memory in the system falls below this value.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.15 Summary

4.3.15.1 Dashboard

The FASTPATH page provides a brief overview of the system and serves as the home page upon successful login to the device.

To access this page, click **System > Summary > Dashboard**.

System Information	
System Description	4 10GBASE X SFP + 16 100/1000BASE T + 8 100/1000BASE X/1000BASE T Combo Managed Industrial Switch, 2.00.005, Linux 3.6.5
System Name	
System Location	
System Contact	
IP Address	192.168.1.167
Burned In MAC Address	74FE48:20:BD:EC
Service Port IP Address	0.0.0.0
Service Port MAC Address	74FE48:20:BD:ED
System Up Time	0 days, 5 hours, 38 mins, 41 secs
Current Time	
Device Information	
Machine Type	4 10GBASE X SFP + 16 100/1000BASE T + 8 100/1000BASE X/1000BASE T Combo Managed Industrial Switch
Machine Model	EKI-9728G
Serial Number	TPAB476058
Software Version	2.00.005
Hardware Version	0x13
Build Version	2017.09.28
Restore Config File Name from USB	
Operating System	Linux 3.6.5
System Resource Usage	
CPU Utilization (60 Second Average)	<div style="width: 15%;"></div> 15 %
Memory Usage	<div style="width: 84%;"></div> 84 %
Disk Space Utilization	
Total Disk Space (Kbytes)	260,096
Free Disk Space (Kbytes)	253,960
Used Disk Space (Kbytes)	6,136

Figure 4.149 System > Summary > Dashboard

The following table describes the items in the previous figure.

Item	Description
System Information	
System Description	The product name of this device.
System Name	The configured name used to identify this device.
System Location	The configured location of this device.
System Contact	The configured contact person for this device.
IP Address	The IP address assigned to the network interface. The network interface is the logical interface that allows remote management of the device via any of the front-panel switch ports.
Burned In MAC Address	The device burned-in universally-administered media access control (MAC) address of the base system.
Service Port IP Address	The IP address assigned to the service port. The service port provides remote management access to the device. Traffic on this port is segregated from operational network traffic on the switch ports and cannot be switched or routed to the operational network.
Service Port MAC Address	The device burned-in universally-administered media access control (MAC) address of the service port.
System Up Time	The time in days, hours, minutes and seconds since the system was last reset.
Current Time	The current time in system.
Device Information	
Machine Type	The device hardware type or product family.
Machine Model	The model identifier, which is usually related to the Machine Type.
Serial Number	The unique device serial number.

Item	Description
FRU Number	The field replaceable unit number.
Maintenance Level	The device hardware change level identifier.
Software Version	The release.version.maintenance number of the software currently running on the device. For example, if the release is 1, the version is 2 and the maintenance number is 4, this version number is displayed as 1.2.4.
Hardware Version	The device hardware version.
Build Version	The software trunk version.
Restore Config File Name from USB	The last time config file name from USB.
Operating System	The device operating system type and version identification information.
System Resource Usage	
CPU Utilization (60 Second Average)	The percentage of CPU utilization for the entire system averaged over the past 60 seconds.
Memory Usage	The percentage of total available system memory (RAM) that is currently in use.
Disk Space Utilization	
Disk Usage	The percentage of total available disk space that is currently in use.
Logged In Users	A brief summary indicating all other users currently logged into the device. The Idle Time field gives an indication of user activity, with a smaller time value denoting more recent access to the system.
Recent Log Entries	A brief list of the newest entries recorded in the system log.
Refresh	Click Refresh to update the screen.

4.3.15.2 Description

Use the System Description page to view and configure basic information about the device. This page contains information that is useful for administrators who manage the device by using a Network Management System (NMS) that communicates with the Simple Network Manage Protocol (SNMP) agent on the device.

To access this page, click **System > Summary > Description**.

Figure 4.150 System > Summary > Description

The following table describes the items in the previous figure.

Item	Description
System Description	The product name of this device.
System Name	The name used to identify this device. The factory default is blank.
System Location	The location of this device. The factory default is blank.
System Contact	The contact person for this device. The factory default is blank.

Item	Description
IP Address	The IP address assigned to the network interface. The network interface is the logical interface that allows remote management of the device via any of the front-panel switch ports.
Service Port IP Address	The IP address assigned to the service port. The service port provides remote management access to the device. Traffic on this port is segregated from operational network traffic on the switch ports and cannot be switched or routed to the operational network.
System Object ID	The base object ID for the device's enterprise MIB. This ID is used for SNMP-based management of the device.
System Up Time	The time in days, hours, minutes, and seconds since the last device reboot.
Current SNTP Synchronized Time	Displays the currently synchronized SNTP time in UTC. If the time is not synchronized with an SNTP server, it displays "Not Synchronized".
MIBs Supported	The list of MIBs supported by the SNMP agent running on this device.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.15.3 Inventory

The System Inventory Information page displays information about the system hardware and software.

To access this page, click **System > Summary > Inventory**.

System Description	4 10GBASE-X SFP + 16 100/1000BASE-T + 8 100/1000BASE-X/1000BASE-T Combo Managed Industrial Switch, 2.00.005, Linux 3.6.5
Machine Type	4 10GBASE-X SFP + 16 100/1000BASE-T + 8 100/1000BASE-X/1000BASE-T Combo Managed Industrial Switch
Machine Model	EKI-9728G
Serial Number	TPAB476058
Manufacturer	Advantech
Burned In MAC Address	74:FE:48:20:BD:EC
Software Version	2.00.005
Hardware Version	0x13
Build Version	2017.09.28
Operating System	Linux 3.6.5
Network Processing Device	BCM56150_A0
Additional Packages	FASTPATH QOS FASTPATH Multicast FASTPATH IPv6 FASTPATH IPv6 Management FASTPATH Routing FASTPATH OpEN API

Refresh

Figure 4.151 System > Summary > Inventory

The following table describes the items in the previous figure.

Item	Description
System Description	The product name of this device.
Machine Type	The hardware platform of this device.
Machine Model	The product model number.
Serial Number	The unique serial number used to identify the device.
Manufacturer	The two-octet code that identifies the manufacturer.
Burned In MAC Address	The device burned-in universally-administered media access control (MAC) address.
Software Version	The release.version.maintenance number of the code currently running on the switch. For example, if the release is 1, the version is 2 and the maintenance number is 4, the format is 1.2.4.
Hardware Version	The device hardware version.
Build Version	The software trunk version.

Item	Description
Operating System	The operating system currently running on the device.
Network Processing Device	Identifies the network processor hardware.
Additional Packages	A list of the optional software packages installed on the device, if any. For example, QoS.
Refresh	Click Refresh to update the screen.

4.3.15.4 MAC Address Table

The MAC address table keeps track of the Media Access Control (MAC) addresses that are associated with each port. This table allows the device to forward unicast traffic through the appropriate port. The MAC address table is sometimes called the bridge table or the forwarding database.

Use the MAC Address Table page to display information about entries in the MAC address table. The transparent bridging function uses these entries to determine how to forward a received frame.

To access this page, click **System > Summary > MAC Address Table**.

VLAN ID	MAC Address	Interface	Interface Index	Status
1	00:08:5B:BD:93:5E	ge0/1	1	Learned
1	00:11:22:33:44:55	CPU	29	Management
1	00:1F:DD:CC:4E:AA	ge0/1	1	Learned
1	00:24:1D:7F:34:04	ge0/1	1	Learned
1	00:26:18:F1:7F:D6	ge0/1	1	Learned
1	00:E0:28:00:00:01	ge0/1	1	Learned
1	1C:6F:65:28:35:44	ge0/1	1	Learned
1	1C:6F:65:28:35:AE	ge0/1	1	Learned
1	1C:6F:65:28:35:B6	ge0/1	1	Learned
1	1C:6F:65:C8:B1:03	ge0/1	1	Learned

Figure 4.152 System > Summary > MAC Address Table

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The VLAN with which the MAC address is associated. A MAC address can be associated with multiple VLANs.
MAC Address	A unicast MAC address for which the switch has forwarding and/or filtering information. The format is a six-byte MAC address, with each byte separated by colons.
Interface	The port where this address was learned. The port identified in this field is the port through which the MAC address can be reached.
Interface Index	The Interface Index of the MIB interface table entry associated with the source port. This value helps identify an interface when using SNMP to manage the device.

Item	Description
Status	<p>Provides information about the entry and why it is in the table, which can be one of the following:</p> <ul style="list-style-type: none"> ■ Static: The address has been manually configured and does not age out. ■ Learned: The address has been automatically learned by the device and can age out when it is not in use. Dynamic addresses are learned by examining information in incoming Ethernet frames. ■ Management: The burned-in MAC address of the device. ■ Self: The MAC address belongs to one of the device's physical interfaces. ■ GMRP Learned: The address was added dynamically by the GARP Multicast Registration Protocol (GMRP). ■ Other: The address was added dynamically through an unidentified protocol or method. ■ Unknown: The device is unable to determine the status of the entry.
Refresh	Click Refresh to update the screen.

4.3.16 Users

4.3.16.1 Accounts

By default, the switch contains two user accounts:

- admin, with 'Read/Write' privilege
- user, with 'Read Only' privileges

Admin account's password is also admin by default. User account's password is also user by default.

If you log on to the switch with the user account that Read/Write privileges (i.e., as admin), you can use the User Accounts page to assign passwords and set security parameters for the default accounts. You can also add up to five read-only accounts. You can delete all accounts except for the Read/Write account.

Note! *Only a user with Read/Write privileges may alter data on this screen, and only one account can exist with Read/Write privileges.*



To access this page, click **System > Users > Accounts**.

User Name	Access Level	Lockout Status	Password Override	Password Expiration	Contained User Group	Operational Permission
admin	Privilege-15	False	Disable		default-usergroup-name	AAA: Read, Write, Execute, Debug OSPF: Read, Write, Execute, Debug
user	Privilege-1	False	Disable			

User Groups

Table is Empty

Refresh Add Edit Remove

Figure 4.153 System > Users > Accounts

The following table describes the items in the previous figure.

Item	Description
User Name	A unique ID or name used to identify this user account.
Access Level	The access or privilege level for this user. The options are: <ul style="list-style-type: none"> ■ Privilege-15: The user can view and modify the configuration. ■ Privilege-1: The user can view the configuration but cannot modify any fields. ■ Privilege-0: The user exists but is not permitted to log on to the device.
Lockout Status	Provides the current lockout status for this user. If the lockout status is True, the user cannot access the management interface even if the correct username and password are provided. The user has been locked out of the system due to a failure to supply the correct password within the configured number of login attempts.
Password Override	Identifies the password override complexity status for this user. <ul style="list-style-type: none"> ■ Enable: The system does not check the strength of the password. ■ Disable: When configuring a password, it is checked against the Strength Check rules configured for passwords.
Password Expiration	Indicates the current expiration date (if any) of the password.
Contained User Group	The associated user groups for the user.
Operational Permission	The operational task permissions for the user.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new user. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

To add a new user:

Click **System > Users > Accounts > Add**.

Figure 4.154 System > Users > Accounts > Add

The following table describes the items in the previous figure.

Item	Description
User Name	A unique ID or name used to identify this user account.
Password	The password assigned to this user.
Confirm	Re-enter the password to confirm that you have entered it correctly.

Item	Description
Access Level	The access or privilege level for this user. The options are: <ul style="list-style-type: none"> ■ Privilege-15: The user can view and modify the configuration. ■ Privilege-1: The user can view the configuration but cannot modify any fields. ■ Privilege-0: The user exists but is not permitted to log on to the device.
Password Override	Identifies the password override complexity status for this user. <ul style="list-style-type: none"> ■ Enable: The system does not check the strength of the password. ■ Disable: When configuring a password, it is checked against the Strength Check rules configured for passwords.
Password Strength	Shows the status of password strength check.
Encrypted Password	Specifies the password encryption.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.3.16.2 Auth Server Users

Use the Auth Server Users page to add and remove users from the local authentication server user database. For some security features, such as IEEE 802.1X port-based authentication, you can configure the device to use the locally stored list of usernames and passwords to provide authentication to users instead of using an external authentication server.

To access this page, click **System > Users > Auth Server Users**.




Figure 4.155 System > Users > Auth Server Users

The following table describes the items in the previous figure.

Item	Description
User Name	A unique name used to identify this user account. You configure the User Name when you add a new user.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new user to the local authentication server database. See the following procedure.
Edit	Click Edit to change the password information for the selected user.
Remove	Click Remove to remove the selected entries.
Clear All User	Click Clear All User to remove all users from the database.

To add a new user:

Click **System > Users > Auth Server Users > Add**.

Figure 4.156 System > Users > Auth Server Users > Add

The following table describes the items in the previous figure.

Item	Description
User Name	A unique name used to identify this user account. You configure the User Name when you add a new user.
Password Required	Select this option to indicate that the user must enter a password to be authenticated. If this option is clear, the user is required only to enter a valid user name.
Password	Specify the password to associate with the user name (if required).
Confirm	Re-enter the password to confirm the entry.
Encrypted	Select this option to encrypt the password before it is stored on the device.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.3.16.3 Sessions

The Logged In Sessions page identifies the users that are logged in to the management interface of the device. The page also provides information about their connections.

To access this page, click **System > Users > Sessions**.

ID	User Name	Connection From	Idle Time	Session Time	Session Type
16	admin	192.168.1.59	00:00:00	02:10:58	HTTP

Figure 4.157 System > Users > Sessions

The following table describes the items in the previous figure.

Item	Description
ID	The unique ID of the session.
User Name	The name that identifies the user account.
Connection From	Identifies the administrative system that is the source of the connection. For remote connections, this field shows the IP address of the administrative system. For local connections through the console port, this field shows the communication standard for the serial connection.
Idle Time	Shows the amount of time in hours, minutes, and seconds that the logged-on user has been inactive.
Session Time	Shows the amount of time in hours, minutes, and seconds since the user logged onto the system.

Item	Description
Session Type	Shows the type of session, which can be Telnet, Serial, SSH, HTTP, or HTTPS.
Refresh	Click Refresh to update the screen.

4.3.16.4 User Domain Name

Use the User Domain Name page to configure the domain name to send to the authentication server, along with the user name and password, to authenticate a user attempting to access the device management interface. Domain name authentication is supported when user authentication is performed by a RADIUS server or TACACS+ server.

To access this page, click **System > Users > User Domain Name**.

ID	User Name	Connection From	Idle Time	Session Time	Session Type
16	admin	192.168.1.59	00:00:00	02:10:58	HTTP

Refresh

Figure 4.158 System > Users > User Domain Name

The following table describes the items in the previous figure.

Item	Description
User Domain Name Mode	The administrative mode of domain name authentication on the device. When enabled, the domain name is included when the user name and password are sent to the authentication server. The domain name can be input by the user in the User Name field on the login screen in a domain-name\username format, or the domain name can be specified in the Domain Name field.
Domain Name	The domain name to send to the authentication server when the user does not provide one in the User Name field during logon. When only the username is provided, the device sends the username as domain-name\username, where domain-name is the string configured in this field. To configure the domain name, click the Edit icon and specify the desired string. To reset the field to its default value, click the Reset icon and confirm the action.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.16.5 Task Groups

The Task Groups page provides the capability to add, edit, and remove task groups.

To access this page, click **System > Users > Task Groups**.





ID	User Name	Connection From	Idle Time	Session Time	Session Type
16	admin	192.168.1.59	00:00:00	02:10:58	HTTP

Refresh

Figure 4.159 System > Users > Task Groups

The following table describes the items in the previous figure.

Item	Description
Task Group	The task group name.
Description	The associated description for task group name.

Item	Description
Parent Task Groups	The associated parent task groups for task group name. To configure this parent task group, click  button in the header row. To remove the parent task group, click  button in the row.
Configured Permission	The configured task permissions for task group.
Operational Permission	The operational task permissions for task group.
Configured Tasks	The list of task names. To configure this task, click  button in the header row. To remove the task, click  button in the row. <ul style="list-style-type: none"> ■ AAA ■ OSPF
Permissions	The task permissions. <ul style="list-style-type: none"> ■ Read ■ Write ■ Debug ■ Execute
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new task group. See the following procedure.
Remove	Click Remove to remove the selected entries.

To add a new task group:

Click **System > Users > Task Groups > Add**.

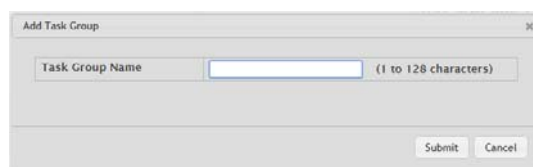


Figure 4.160 System > Users > Task Groups > Add

The following table describes the items in the previous figure.

Item	Description
Task Group Name	Enter the name of task group.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.3.16.6 User Groups

The User Group page provides the capability to add, edit, and remove user groups.

To access this page, click **System > Users > User Groups**.



Figure 4.161 System > Users > User Groups

The following table describes the items in the previous figure.

Item	Description
User Group	The user group name.

Item	Description
Description	The associated description for user group name.
Parent User Groups	The associated parent user groups for user group. To configure this parent user group, click + button in the header row. To remove the parent user group, click ⏻ button in the row.
Contained Task Group	The associated task groups for user group. To configure this task group, click + button in the header row. To remove the task group, click ⏻ button in the row.
Operational Permission	The operational task permissions for user group.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new user group. See the following procedure.
Remove	Click Remove to remove the selected entries.

To add a new user group:

Click **System > Users > User Groups > Add**.




Figure 4.162 System > Users > User Groups > Add

The following table describes the items in the previous figure.

Item	Description
User Group Name	Enter the name of user group.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.3.17 Utilities

4.3.17.1 System Reset

Use the System Reset page to reboot the system. If the platform supports stacking, you can reset any of the switches in the stack, or all switches in the stack from this page.

To access this page, click **System > Utilities > System Reset**.




Figure 4.163 System > Utilities > System Reset

The following table describes the items in the previous figure.

Item	Description
Generate Core Dump before reset	Generates core dump file on demand.
Reset	Click Reset to initiates the system reset action after displaying a confirmation message.

4.3.17.2 Ping

Use the Ping page to tell the switch to send a Ping request to a specified IP address. You can use this feature to check whether the switch can communicate with a particular network host.

To access this page, click **System > Utilities > Ping**.

Figure 4.164 System > Utilities > Ping

The following table describes the items in the previous figure.

Item	Description
Host Name or IP Address	The DNS-resolvable hostname or IP address of the system to ping.
Count	Enter the number of ICMP echo request packets to send to the host.
Interval (Seconds)	Enter the number of seconds to wait between sending ping packets.
Size (Bytes)	The size of the ping packet, in bytes. Changing the size allows you to troubleshoot connectivity issues with a variety of packet sizes, such as large or very large packets.
Source	The source IP address or interface to use when sending the echo request packets. If source is not required, select None as source option.
IP Address	The source IP address to use when sending the Echo requests packets. This field is enabled when IP Address is selected as source option.
Interface	The interface to use when sending the Echo requests packets. This field is enabled when Interface is selected as source option.
Status	The current status of the ping test, which can be: <ul style="list-style-type: none"> ■ Not Started: The ping test has not been initiated since viewing the page. ■ In Progress: The ping test has been initiated and is running. ■ Stopped: The ping test was interrupted by clicking the Stop button. ■ Done: The test has completed, and information about the test is displayed in the Results area.
Results	The results of the ping test, which includes information about the reply (if any) received from the host.
Start	Click Start to start the ping test. The device sends the specified number of ping packets to the host.
Stop	Click Stop to interrupts the current ping test.

4.3.17.3 Ping IPv6

Use the Ping IPv6 page to tell the device to send one or more ping requests to a specified IPv6 host. You can use the ping request to check whether the device can communicate with a particular host on an IPv6 network. A ping request is an Internet Control Message Protocol version 6 (ICMPv6) echo request packet. The information you enter on this page is not saved as part of the device configuration.

To access this page, click **System > Utilities > Ping IPv6**.

Figure 4.165 System > Utilities > Ping IPv6

The following table describes the items in the previous figure.

Item	Description
Ping	Select either a global IPv6 address or a link local address to ping. A global address is routable over the Internet, while a link-local address is intended for communication only within the local network. Link local addresses have a prefix of fe80::/64.
Interface	Select the interface on which to issue the Link Local ping request.
Host Name or IPv6 Address	Enter the global or link-local IPv6 address, or the DNS-resolvable host name of the station to ping. If the ping type is Link Local, you must enter a link-local address and cannot enter a host name.
Count	Enter the number of ICMP echo request packets to send to the host.
Interval (Seconds)	Enter the number of seconds to wait between sending ping packets.
Size (Bytes)	The size of the ping packet, in bytes. Changing the size allows you to troubleshoot connectivity issues with a variety of packet sizes, such as large or very large packets.
Source	The source IP address or interface to use when sending the echo request packets. If source is not required, select None as source option.
IPv6 Address	The source IPv6 address to use when sending the Echo requests packets. This field is enabled when IP Address is selected as source option.
Interface	The interface to use when sending the Echo requests packets. This field is enabled when Interface is selected as source option.
Results	The results of the ping test, which includes information about the reply (if any) received from the host.
Submit	Click Submit to start the ping test.

4.3.17.4 TraceRoute

Use the TraceRoute page to determine the layer 3 path a packet takes from the device to a specific IP address or hostname. When you initiate the TraceRoute command by clicking the **Start** button, the device sends a series of TraceRoute probes toward the destination. The results list the IP address of each layer 3 device a probe passes through until it reaches its destination - or fails to reach its destination and is discarded. The information you enter on this page is not saved as part of the device configuration.

To access this page, click **System > Utilities > TraceRoute**.

Figure 4.166 System > Utilities > TraceRoute

The following table describes the items in the previous figure.

Item	Description
Host Name or IP Address	The DNS-resolvable hostname or IP address of the system to attempt to reach.
Probes Per Hop	TraceRoute works by sending UDP packets with increasing Time-To-Live (TTL) values. Specify the number of probes sent with each TTL.
MaxTTL	The maximum Time-To-Live (TTL). The TraceRoute terminates after sending probes that can be layer 3 forwarded this number of times. If the destination is further away, the TraceRoute will not reach it.
InitTTL	The initial Time-To-Live (TTL). This value controls the maximum number of layer 3 hops that the first set of probes may travel.
MaxFail	The number of consecutive failures that terminate the TraceRoute. If the device fails to receive a response for this number of consecutive probes, the TraceRoute terminates.
Interval (Seconds)	The number of Seconds to wait between sending probes.
Port	The UDP destination port number to be used in probe packets. The port number should be a port that the target host is not listening on, so that when the probe reaches the destination, it responds with an ICMP Port Unreachable message.
Size (Bytes)	The size of probe payload in bytes.
Source	Select None, IP Address, Interface, or Loopback as a source. Each option enables the respective menu item: IP Address, Interface, or Interface Loopback, allowing the entry of the related information.
IP Address	Enabled if IP Address option is selected from Source setting.
Interface	Enabled if Interface option is selected from Source setting.
Interface Loopback	Enabled if Loopback option is selected from Source setting.

Item	Description
Status	<p>The current status of the TraceRoute, which can be:</p> <ul style="list-style-type: none"> ■ Not Started: The TraceRoute has not been initiated since viewing the page. ■ In Progress: The TraceRoute has been initiated and is running. ■ Stopped: The TraceRoute was interrupted by clicking the Stop button. ■ Done: The TraceRoute has completed, and information about the TraceRoute is displayed in the Results area.
Results	<p>The results of the TraceRoute, which are displayed in the following format:</p> <pre> 1 10.20.24.1 0 ms 0 ms 0 ms 2 66.20.17.9 10 ms 0 ms 10 ms 3 66.20.246.82 10 ms 20 ms 10 ms 4 129.20.4.4 20 ms 10 ms 40 ms 5 129.20.3.55 80 ms 80 ms 90 ms 6 129.20.5.246 80 ms 80 ms 80 ms 7 198.20.90.26 70 ms 70 ms 70 ms 8 216.20.255.105 90 ms 70 ms 80 ms 9 63.20.216.155 80 ms 80 ms 90 ms </pre> <p>Hop Count = 9 Last TTL = 9 Test attempt = 27 Test Success = 27</p> <p>For each TTL value probed, the results show the IP address of the router that responded to the probes and the response time for each probe. If no response is received for probes with a particular TTL, the IP address is reported as 0.0.0.0.</p> <p>An error code may be printed with the response time for each probe. The error codes signify that either no response was received or an ICMP Destination Unreachable message was received with error codes as follows:</p> <ul style="list-style-type: none"> ■ * no response was received to the probe ■ P: Protocol unreachable (RFC 792) ■ N: Network unreachable (RFC 792) ■ H: Host unreachable (RFC 792) ■ F: Fragmentation needed and DF set (RFC 792) ■ S: Source route failed (RFC 792) ■ A: Communication with Destination Network is Administratively Prohibited (RFC 1122) ■ C: Communication with Destination Host is Administratively Prohibited (RFC 1122) <p>The Hop Count is the number of sets of probes sent, each set of probes having a particular TTL. The Last TTL is the TTL sent in the final set of probes. The Test Attempt value shows the number of probes sent. The Test Success value shows the number of probes that received a response.</p>
Start	Click Start to initiates the TraceRoute.
Stop	Click Stop to interrupts the running TraceRoute.

4.3.17.5 TraceRoute IPv6

Use the TraceRoute IPv6 page to determine the layer 3 path a packet takes from the device to a specific IP address or hostname. When you initiate the IPv6 TraceRoute command by clicking the **Submit** button, the device sends a series of IPv6 TraceRoute probes toward the destination. The results list the IP address of each layer 3 device a probe passes through until it reaches its destination - or fails to reach its destination and is discarded. The information you enter on this page is not saved as part of the device configuration.

To access this page, click **System > Utilities > TraceRoute IPv6**.

Figure 4.167 System > Utilities > TraceRoute IPv6

The following table describes the items in the previous figure.

Item	Description
Host Name or IPv6 Address	The DNS-resolvable hostname or IPv6 address of the system to attempt to reach.
Probes Per Hop	IPv6 TraceRoute works by sending UDP packets with increasing Time-To-Live (TTL) values. Specify the number of probes sent with each TTL.
MaxTTL	The maximum Time-To-Live (TTL). The TraceRoute terminates after sending probes that can be layer 3 forwarded this number of times. If the destination is further away, the TraceRoute will not reach it.
InitTTL	The initial Time-To-Live (TTL). This value controls the maximum number of layer 3 hops that the first set of probes may travel.
MaxFail	The number of consecutive failures that terminate the TraceRoute. If the device fails to receive a response for this number of consecutive probes, the TraceRoute terminates.
Interval (Seconds)	Specifies the time between probes, in Seconds. If a response is not received within this interval, then traceroute considers the probe a failure and sends the next probe. If traceroute does receive a response to a probe within this interval, then it sends the next probe immediately.
Port	The UDP destination port number to be used in probe packets. The port number should be a port that the target host is not listening on, so that when the probe reaches the destination, it responds with an ICMPv6 Port Unreachable message.
Size (Bytes)	The size of probe payload in bytes.
Source	The source IP address or interface to use when sending the trace route command. If source is not required, select None as source option.

Item	Description
IPv6 Address	The source IPv6 address to use when sending the trace route command. This field is enabled when IP Address is selected as source option.
Interface	The interface to use when sending the trace route command. This field is enabled when Interface is selected as source option.
Results	<p>The results of the TraceRoute, which are displayed in the following format:</p> <pre> 1 3001::1 708 ms 41 ms 11 ms 2 4001::2 250 ms 200 ms 193 ms 3 5001::3 289 ms 313 ms 278 ms 4 6001::4 651 ms 41 ms 270 ms 5 :: * N * N * N </pre> <p>Hop Count = 4 Last TTL = 5 Test attempt = 1 Test Success = 0 For each TTL value probed, the results show the IP address of the router that responded to the probes and the response time for each probe. If no response is received for probes with a particular TTL, the IP address is reported as 0.0.0.0.</p> <p>An error code may be printed with the response time for each probe. The error codes signify that either no response was received or an ICMP Destination Unreachable message was received with error codes as follows:</p> <ul style="list-style-type: none"> ■ * no response was received to the probe ■ P: Protocol unreachable (RFC 792) ■ N: Network unreachable (RFC 792) ■ H: Host unreachable (RFC 792) ■ F: Fragmentation needed and DF set (RFC 792) ■ S: Source route failed (RFC 792) ■ A: Communication with Destination Network is Administratively Prohibited (RFC 1122) ■ C: Communication with Destination Host is Administratively Prohibited (RFC 1122) <p>The Hop Count is the number of sets of probes sent, each set of probes having a particular TTL. The Last TTL is the TTL sent in the final set of probes. The Test Attempt value shows the number of probes sent. The Test Success value shows the number of probes that received a response.</p>
Submit	Click Submit to initiates the TraceRoute.

4.3.17.6 IP Address Conflict

Use the IP Address Conflict Detection page to determine whether the IP address configured on the device is the same as the IP address of another device on the same LAN (or on the Internet, for a routable IP address) and to help you resolve any existing conflicts. An IP address conflict can make both this system and the system with the same IP address unusable for network operation.

To access this page, click **System > Utilities > IP Address Conflict**.

Status	
IP Address Conflict Currently Exists	False
History	
Last Conflicting IP Address	
Last Conflicting MAC Address	
Time Since Conflict Detected	
<input type="button" value="Refresh"/> <input type="button" value="Run Detection"/> <input type="button" value="Clear History"/>	

Figure 4.168 System > Utilities > IP Address Conflict

The following table describes the items in the previous figure.

Item	Description
Status	
IP Address Conflict Currently Exists	<p>Indicates whether a conflicting IP address has been detected since this status was last reset.</p> <ul style="list-style-type: none"> ■ False: No conflict detected (the subsequent fields on this page display as N/A). ■ True: Conflict was detected (the subsequent fields on this page show the relevant information).
History	
Last Conflicting IP Address	The device interface IP address that is in conflict. If multiple conflicts were detected, only the most recent occurrence is displayed.
Last Conflicting MAC Address	The MAC address of the remote host associated with the IP address that is in conflict. If multiple conflicts are detected, only the most recent occurrence is displayed.
Time Since Conflict Detected	The elapsed time (displayed in days, hours, minutes, and seconds) since the last address conflict was detected, provided the Clear History button has not yet been pressed.
Refresh	Click Refresh to update the screen.
Run Detection	Click Run Detection to activate the IP address conflict detection operation in the system.
Clear History	Click Clear History to reset the IP address conflict detection status information that was last seen by the device.

4.3.17.7 Transfer

Use the File Transfer page to upload files from the device to a remote system and to download files from a remote system to the device.


To access this page, click **System > Utilities > Transfer**.


Transfer Protocol	Upload <i>Transfer a file from the device</i>	Download <i>Transfer a file to the device</i>
HTTP		
TFTP		
FTP		


Figure 4.169 System > Utilities > Transfer

The following table describes the items in the previous figure.

Item	Description
Transfer Protocol	The protocol to use to transfer the file. Files can be transferred from the device to a remote system using TFTP, or FTP. Files can be transferred from a remote system to the device using HTTP, TFTP, or FTP.
Upload	To transfer a file from the device to a remote system using TFTP, or FTP, click the upload icon in the same row as the desired transfer protocol. The File Upload window appears. Configure the information for the file transfer (described below), and click the upload icon to the right of the Progress field to begin the transfer.
Download	To transfer a file from a remote system to the device using HTTP, TFTP, or FTP, click the download icon in the same row as the desired transfer protocol. The File Download window appears. Configure the information for the file transfer (described below), and click the download icon to the right of the Progress field to begin the transfer.

Item	Description
	<ul style="list-style-type: none"> ■ File Type: Specify the type of file to transfer from the device to a remote system. <ul style="list-style-type: none"> – Active Code: Select this option to transfer an active image. – Backup Code: Select this option to transfer an backup image. – Startup Configuration: Select this option to transfer a copy of the stored startup configuration from the device to a remote system. – Backup Configuration: Select this option to transfer a copy of the stored backup configuration from the device to a remote system. – Script File: Select this option to transfer a custom text configuration script from the device to a remote system. – CLI Banner: Select this option to transfer the file containing the text to be displayed on the CLI before the login prompt to a remote system. – MIB File: Select this option to transfer the MIB file to a remote system. – Crash Log: Select this option to transfer the system crash log to a remote system. – Operational Log: Select this option to transfer the system operational log to a remote system. – Startup Log: Select this option to transfer the system startup log to a remote system. – Trap Log: Select this option to transfer the system trap records to a remote system. – Error Log: Select this option to transfer the system error (persistent) log, which is also known as the event log, to a remote system. – Buffered Log: Select this option to transfer the system buffered (in-memory) log to a remote system. ■ Image: If the selected File Type is Code, specify whether to transfer the Active or Backup image to a remote system. ■ Server Address: Specify the IPv4 address, IPv6 address, or DNS-resolvable hostname of the remote server that will receive the file. ■ File Path: Specify the path on the server where you want to put the file. ■ File Name: Specify the name that the file will have on the remote server. ■ User Name: For and FTP transfers, if the server requires authentication, specify the user name for remote login to the server that will receive the file. ■ Password: For and FTP transfers, if the server requires authentication, specify the password for remote login to the server that will receive the file. ■ Progress: Represents the completion percentage of the file transfer. The file transfer begins after you complete the required fields and click the upload icon to the right of this field. ■ Status: Provides information about the status of the file transfer.

Item	Description
	<ul style="list-style-type: none"> ■ File Type: Specify the type of file to transfer to the device: <ul style="list-style-type: none"> – Active Code: Select this option to transfer a new image to the device. The code file is stored as the active image. – Backup Code: Select this option to transfer a new image to the device. The code file is stored as the backup image. – Startup Configuration: Select this option to update the stored startup configuration file. If the file has errors, the update will be stopped. – Backup Configuration: Select this option to update the stored backup configuration file. If the file has errors, the update will be stopped. – Script File: Select this option to transfer a text-based configuration script to the device. You must use the command-line interface (CLI) to validate and activate the script. – CLI Banner: Select this option to transfer the CLI banner file to the device. This file contains the text to be displayed on the CLI before the login prompt. – IAS Users: Select this option to transfer an Internal Authentication Server (IAS) users database file to the device. The IAS user database stores a list of user name and (optional) password values for local port-based user authentication. – SSH-1 RSA Key File: Select this option to transfer an SSH-1 Rivest-Shamir-Adleman (RSA) key file to the device. SSH key files contain information to authenticate SSH sessions for remote CLI-based access to the device. – SSH-2 RSA Key PEM File: Select this option to transfer an SSH-2 Rivest-Shamir-Adleman (RSA) key file (PEM Encoded) to the device. – SSH-2 DSA Key PEM File: Select this option to transfer an SSH-2 Digital Signature Algorithm (DSA) key file (PEM Encoded) to the device. – SSL Trusted Root Certificate PEM File: Select this option to transfer an SSL Trusted Root Certificate file (PEM Encoded) to the device. SSL files contain information to encrypt, authenticate, and validate HTTPS sessions. – SSL Server Certificate PEM File: Select this option to transfer an SSL Server Certificate file (PEM Encoded) to the device. – SSL DH Weak Encryption Parameter PEM File: Select this option to transfer an SSL Diffie-Hellman Weak Encryption Parameter file (PEM Encoded) to the device. – SSL DH Strong Encryption Parameter PEM File: Select this option to transfer an SSL Diffie-Hellman Strong Encryption Parameter file (PEM Encoded) to the device. <p><i>Note: To download SSH key files, SSH must be administratively disabled, and there can be no active SSH sessions. To download SSL related files, HTTPS must be administratively disabled.</i></p>

Item	Description
	<ul style="list-style-type: none"> ■ Select File: If HTTP is the Transfer Protocol, browse to the directory where the file is located and select the file to transfer to the device. This field is not present if the Transfer Protocol is TFTP or FTP. ■ Server Address: For TFTP, or FTP transfers, specify the IPv4 address, IPv6 address, or DNS-resolvable hostname of the remote server. ■ File Path: For TFTP, or FTP transfers, specify the path on the server where the file is located. ■ File Name: For TFTP, or FTP transfers, specify the name of the file you want to transfer to the device. ■ User Name: For FTP transfers, if the server requires authentication, specify the user name for remote login to the server where the file resides. ■ Password: For FTP transfers, if the server requires authentication, specify the password for remote login to the server where the file resides. ■ Progress: Represents the completion percentage of the file transfer. The file transfer begins after you complete the required fields and click the download icon to the right of this field. ■ Status: Provides information about the status of the file transfer.

4.3.17.8 Digital Signature Verification

Use the Digital Signature Verification page to configure digital signature verification on downloading files from a remote system to the device.

To access this page, click **System > Utilities > Digital Signature Verification**.




Figure 4.170 System > Utilities > Digital Signature Verification

The following table describes the items in the previous figure.

Item	Description
Digital Signature Verification	Provides option to verify the digital signature of a downloaded file.
Code	Verify the digital signature of downloaded code image files.
Configuration	Verify the digital signature of downloaded configuration script files.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.3.17.9 Core Dump

Use the Core Dump page to configure Core Dump feature.

To access this page, click **System > Utilities > Core Dump**.

Figure 4.171 System > Utilities > Core Dump

The following table describes the items in the previous figure.

Item	Description
Core Dump Configuration	
Protocol	The protocol used to store the core dump file. User can select one of the options to configure: <ul style="list-style-type: none"> None: Disable Core Dump. TFTP: Configure protocol to upload Core Dump to the TFTP server. NFS: Configure protocol to upload Core Dump to the NFS share. USB: Configure protocol to upload Core Dump to the USB mount point. FTP: Configure protocol to upload Core Dump to the FTP server.
Core Dump File Name Prefix	Prefix for the Core Dump file name. If hostname is configured, it takes else while generating Core Dump file. The prefix length is 15 characters.
Use Host Name	To use hostname (or MAC if hostname is not configured) to name Core Dump file.
Use Time Stamp	To use timestamp to name Core Dump file.
TFTP IP Address	IP address of remote TFTP server to dump core file to external server.
FTP IP Address	IP address of remote FTP server to dump core file to external server.
FTP Username	Username of remote FTP server.
FTP Password	Password of remote FTP server.
File Path	File-path to dump core file to TFTP server, NFS mount or USB device sub-directory.
Switch Chip Registers Dump	To enable or disable switch-chip-register dump in case of an exception. The switch-chip-register dump is taken only for master unit and not for member units.
Stack IP Address Protocol	Protocol (DHCP or Static) to be used to configure service port when a unit has crashed. If configured as DHCP then the unit gets the IP address from DHCP server available in the network. If configured as Static, an IP address from the Core Dump Stack IP Address Pool is used.

Item	Description
Core Dump Stack IP Address Pool	
IP Address	Static IP address to be assigned to individual units service port in the stack when the switch has crashed. This IP address is used to perform the core dump.
Host Mask	The subnet mask.
Default Router Address	The IP address of the router.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new stack IP address.
Remove	Click Remove to remove the selected entries.

4.3.17.10 Core Dump Test

Use the Core Dump Test page to test the core dump setup. For example if protocol is configured as TFTP, it communicates with TFTP server and informs user if the TFTP server can be contacted.

To access this page, click **System > Utilities > Core Dump Test**.



Figure 4.172 System > Utilities > Core Dump Test

The following table describes the items in the previous figure.

Item	Description
Status	Displays test status as Ok if test passes and Error if test fails.
Result	Displays detailed error information with logs.
Submit	Click Submit to save the values and update the screen.

4.4 Switching

4.4.1 Auto Recovery

4.4.1.1 Configuration

Use the Configuration page to set the global Auto Recovery options for Components. To access this page, click **Switching > Auto Recovery > Configuration**.

Figure 4.173 Switching > Auto Recovery > Configuration

The following table describes the items in the previous figure.

Item	Description
Auto Recovery Components	<p>This enables/disables auto recovery for the specified component (e.g. BPDUGuard). An interface in the diagnostic disabled state for the configured components is recovered (link up) when the recovery interval expires. If the interface continues to encounter errors (from any listed components), it may be placed back in the diagnostic disabled state and the interface will be disabled (link down). Interfaces in the diagnostic disabled state may also be manually recovered by enabling them in Port Summary page.</p> <ul style="list-style-type: none"> ■ ARP Inspection ■ BPDUGuard ■ BPDURateLimit ■ BroadcastStormControl ■ DenialOfService ■ DHCPRateLimit ■ Keepalive ■ MACLocking ■ MulticastStormControl ■ UDLD ■ UnicastStormControl
Auto Recovery Parameters	
Recovery Time	This configures the auto recovery time interval. The auto recovery time interval is common for all components. The default value of the timer is 300 seconds and the range is from 30 to 86400.
D-Disabled Interface Status	
Interface	The interface which is error disabled.
Admin Mode	The administrative mode of the interface.

Item	Description
Port Status	Indicates whether the link is up or down. The link is the physical connection between the port or trunk and the interface on another device.
Error Disable Reason	<p>If the device detects an error condition for an interface, then the device puts the interface in error disabled state by placing the interface in diagnostic disabled state. Following are the reasons due to which the interface can go into error disable state.</p> <ul style="list-style-type: none"> ■ ARP Inspection ■ BPDU Guard ■ BPDU Storm ■ Broadcast Storm ■ Denial Of Service ■ DHCP Rate Limit ■ Keepalive ■ MAC Locking ■ Multicast Storm ■ UDLD ■ Unicast Storm
Auto Recovery Time Left (Seconds)	When Auto Recovery is enabled and the interface is placed in diagnostic disabled state, then a recovery timer starts for that interface. Once this timer expires, the device checks if the interface is in diagnostic disabled state. If yes, then the device enables the diagnostic disabled interface.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.4.2 Class of Service

4.4.2.1 802.1p

The IEEE 802.1p feature allows traffic prioritization at the MAC level. The switch can prioritize traffic based on the 802.1p tag attached to the L2 frame. Each port on the switch has multiple queues to give preference to certain packets over others based on the class of service (CoS) criteria you specify. When a packet is queued for transmission in a port, the rate at which it is serviced depends on how the queue is configured and possibly the amount of traffic present in the other queues of the port. If a delay is necessary, packets get held in the queue until the scheduler authorizes the queue for transmission.

Use the 802.1p Priority Mapping page to view or change which internal traffic classes are mapped to the 802.1p priority class values in Ethernet frames the device receives. The priority-to-traffic class mappings can be applied globally or per-interface. The mapping allows the device to group various traffic types (e.g. data or voice) based on their latency requirements and give preference to time-sensitive traffic.

To access this page, click **Switching > Class of Service > 802.1p**.

Interface	Priority 0	Priority 1	Priority 2	Priority 3	Priority 4	Priority 5	Priority 6	Priority 7
<input type="checkbox"/> Global	1	0	0	1	2	2	3	3
<input type="checkbox"/> ge0/1	1	0	0	1	2	2	3	3
<input type="checkbox"/> ge0/2	1	0	0	1	2	2	3	3
<input type="checkbox"/> ge0/3	1	0	0	1	2	2	3	3
<input type="checkbox"/> ge0/4	1	0	0	1	2	2	3	3
<input type="checkbox"/> ge0/5	1	0	0	1	2	2	3	3
<input type="checkbox"/> ge0/6	1	0	0	1	2	2	3	3
<input type="checkbox"/> ge0/7	1	0	0	1	2	2	3	3
<input type="checkbox"/> ge0/8	1	0	0	1	2	2	3	3
<input type="checkbox"/> ge0/9	1	0	0	1	2	2	3	3

Figure 4.174 Switching > Class of Service > 802.1p

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. The Global entry represents the common settings for all interfaces, unless specifically overridden individually.
Priority	The heading row lists each 802.1p priority value (0-7), and the data in the table shows which traffic class is mapped to the priority value. Incoming frames containing the designated 802.1p priority value are mapped to the corresponding traffic class in the device.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.

4.4.3 DHCP Snooping

4.4.3.1 Base

Global

Use the DHCP Snooping Configuration page to view and configure the global settings for DHCP Snooping.

To access this page, click **Switching > DHCP Snooping > Base > Global**.

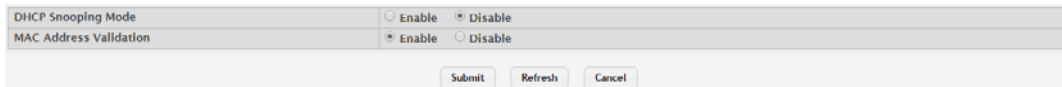


Figure 4.175 Switching > DHCP Snooping > Base > Global

The following table describes the items in the previous figure.

Item	Description
DHCP Snooping Mode	The administrative mode of DHCP snooping on the device.
MAC Address Validation	Enables or Disables the verification of the sender MAC address for DHCP snooping. When enabled, the device checks packets that are received on untrusted interfaces to verify that the MAC address and the DHCP client hardware address match. If the addresses do not match, the device drops the packet.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

VLAN Configuration

Use the DHCP Snooping VLAN Configuration page to view and configure the DHCP snooping settings on VLANs that exist on the device. DHCP snooping can be configured on switching VLANs and routing VLANs. For Layer 2 (non-routing) VLANs, DHCP snooping forwards valid DHCP client messages received on the VLANs. The message is forwarded on all trusted interfaces in the VLAN. When a DHCP packet is received on a routing VLAN, the DHCP snooping application applies its filtering rules and updates the bindings database. If a client message passes filtering rules, the message is placed into the software forwarding path, where it may be processed by the DHCP relay agent, the local DHCP server, or forwarded as an IP packet.

To access this page, click **Switching > DHCP Snooping > Base > VLAN Configuration**.

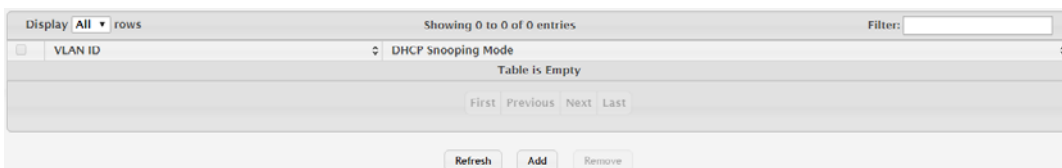


Figure 4.176 Switching > DHCP Snooping > Base > VLAN Configuration

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The VLAN ID that is enabled for DHCP snooping. In the Add DHCP Snooping VLAN Configuration window, this field lists the VLAN ID of all VLANs that exist on the device.
DHCP Snooping Mode	The current administrative mode of DHCP snooping for the VLAN. Only VLANs that are enabled for DHCP snooping appear in the list.

Item	Description
Refresh	Click Refresh to update the screen.
Add	Click Add to enable a VLAN for DHCP snooping. To select multiple VLANs, CTRL + click each VLAN to select. See the following procedure.
Remove	Click Remove to disable DHCP snooping on the selected entries.

To enable a VLAN for DHCP snooping:

Click **Switching > DHCP Snooping > Base > VLAN Configuration > Add**.

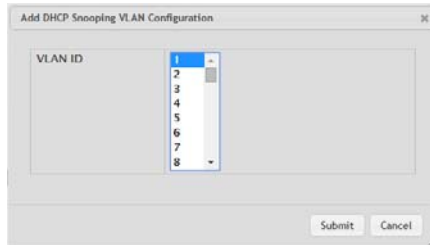


Figure 4.177 Switching > DHCP Snooping > Base > VLAN Configuration > Add

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The VLAN ID that is enabled for DHCP snooping. In the Add DHCP Snooping VLAN Configuration window, this field lists the VLAN ID of all VLANs that exist on the device.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

Interface Configuration

Use the DHCP Snooping Interface Configuration page to view and configure the DHCP snooping settings for each interface. The DHCP snooping feature processes incoming DHCP messages. For DHCPRELEASE and DHCPDECLINE messages, the feature compares the receive interface and VLAN with the client's interface and VLAN in the binding database. If the interfaces do not match, the application logs the event (when logging of invalid packets is enabled) and drops the message. If MAC address validation is globally enabled, messages that pass the initial validation are checked to verify that the source MAC address and the DHCP client hardware address match. Where there is a mismatch, DHCP snooping logs the event (when logging of invalid packets is enabled) and drops the packet.

To access this page, click **Switching > DHCP Snooping > Base > Interface Configuration**.

Interface	Trust State	Log Invalid Packets	Rate Limit (pps)	Burst Interval (Seconds)
ge0/1	Disabled	Disabled		
ge0/2	Disabled	Disabled		
ge0/3	Disabled	Disabled		
ge0/4	Disabled	Disabled		
ge0/5	Disabled	Disabled		
ge0/6	Disabled	Disabled		
ge0/7	Disabled	Disabled		
ge0/8	Disabled	Disabled		
ge0/9	Disabled	Disabled		
ge0/10	Disabled	Disabled		

Figure 4.178 Switching > DHCP Snooping > Base > Interface Configuration

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. When configuring the settings for one or more interfaces, this field identifies each interface that is being configured.
Trust State	The trust state configured on the interface. The trust state is one of the following: <ul style="list-style-type: none"> ■ Disabled: The interface is considered to be untrusted and could potentially be used to launch a network attack. DHCP server messages are checked against the bindings database. On untrusted ports, DHCP snooping enforces the following security rules: <ul style="list-style-type: none"> – DHCP packets from a DHCP server (DHCPOFFER, DHCPACK, DHCPNAK, DHCPRELEASEQUERY) are dropped. – DHCPRELEASE and DHCPDECLINE messages are dropped if the MAC address is in the snooping database but the binding's interface is other than the interface where the message was received. – DHCP packets are dropped when the source MAC address does not match the client hardware address if MAC Address Validation is globally enabled. ■ Enabled: The interface is considered to be trusted and forwards DHCP server messages without validation.
Log Invalid Packets	The administrative mode of invalid packet logging on the interface. When enabled, the DHCP snooping feature generates a log message when an invalid packet is received and dropped by the interface.
Rate Limit (pps)	The rate limit value for DHCP packets received on the interface. To prevent DHCP packets from being used as a DoS attack when DHCP snooping is enabled, the snooping application enforces a rate limit for DHCP packets received on untrusted interfaces. If the incoming rate of DHCP packets exceeds the value of this object during the amount of time specified for the burst interval, the port will be shutdown. You must administratively enable the port to allow it to resume traffic forwarding.
Burst Interval (Seconds)	The burst interval value for rate limiting on this interface. If the rate limit is unspecified, then burst interval has no meaning.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.

Static Bindings

Use the DHCP Snooping Static Bindings page to view, add, and remove static bindings in the DHCP snooping bindings database.

To access this page, click **Switching > DHCP Snooping > Base > Static Bindings**.




Figure 4.179 Switching > DHCP Snooping > Base > Static Bindings

The following table describes the items in the previous figure.

Item	Description
Interface	The interface on which the DHCP client is authorized.
MAC Address	The MAC address associated with the DHCP client. This is the Key to the binding database.
VLAN ID	The ID of the VLAN the client is authorized to use.
IP Address	The IP address of the client.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a static entry to the DHCP snooping bindings table. See the following procedure.
Remove	Click Remove to remove the selected entries.

To add a static entry to the DHCP snooping bindings table:

Click **Switching > DHCP Snooping > Base > Static Bindings > Add**.




Figure 4.180 Switching > DHCP Snooping > Base > Static Bindings > Add

The following table describes the items in the previous figure.

Item	Description
Interface	The interface on which the DHCP client is authorized.
MAC Address	The MAC address associated with the DHCP client. This is the Key to the binding database.
VLAN ID	The ID of the VLAN the client is authorized to use.
IP Address	The IP address of the client.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

Dynamic Bindings

Use the DHCP Snooping Dynamic Bindings page to view and clear dynamic bindings in the DHCP snooping bindings database. The DHCP snooping feature uses DHCP messages to build and maintain the bindings database. The bindings database includes data for clients only on untrusted ports. DHCP snooping creates a tentative binding from DHCP DISCOVER and REQUEST messages. Tentative bindings tie a client to an interface (the interface where the DHCP client message was received). Tentative bindings are completed when DHCP snooping learns the client's IP address from a DHCP ACK message on a trusted port. DHCP snooping removes bindings in response to DECLINE, RELEASE, and NACK messages. The DHCP snooping feature ignores the ACK messages as a reply to the DHCP Inform messages received on trusted ports.

To access this page, click **Switching > DHCP Snooping > Base > Dynamic Bindings**.

Figure 4.181 Switching > DHCP Snooping > Base > Dynamic Bindings

The following table describes the items in the previous figure.

Item	Description
Interface	The interface on which the DHCP client message was received.
MAC Address	The MAC address associated with the DHCP client that sent the message. This is the Key to the binding database.
VLAN ID	The VLAN ID of the client interface.
IP Address	The IP address assigned to the client by the DHCP server.
Lease Time	The remaining IP address lease time for the client.
Refresh	Click Refresh to update the screen.
Clear	Click Clear to remove the selected entries in the database.

Persistent

Use the DHCP Snooping Persistent Configuration page to configure the persistent location of the DHCP snooping bindings database. The bindings database can be stored locally on the device or on a remote system somewhere else in the network. The device must be able to reach the IP address of the remote system to send bindings to a remote database.

To access this page, click **Switching > DHCP Snooping > Base > Persistent**.

Figure 4.182 Switching > DHCP Snooping > Base > Persistent

The following table describes the items in the previous figure.

Item	Description
Store	The location of the DHCP snooping bindings database, which is either locally on the device (Local) or on a remote system (Remote).

Item	Description
Remote IP Address	The IP address of the system on which the DHCP snooping bindings database will be stored. This field is available only if Remote is selected in the Store field.
Remote File Name	The file name of the DHCP snooping bindings database in which the bindings are stored. This field is available only if Remote is selected in the Store field.
Write Delay (Seconds)	The amount of time to wait between writing bindings information to persistent storage. This allows the device to collect as many entries as possible (new and removed) before writing them to the persistent file.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

Statistics

Use the DHCP Snooping Statistics page to view and clear per-interface statistics about the DHCP messages filtered by the DHCP snooping feature. Only interfaces that are enabled for DHCP snooping and are untrusted appear in the table.

To access this page, click **Switching > DHCP Snooping > Base > Statistics**.

Interface	MAC Verify Failures	Client Ifc Mismatch	DHCP Server Msgs Received
ge0/1	0	0	0
ge0/2	0	0	0
ge0/3	0	0	0
ge0/4	0	0	0
ge0/5	0	0	0
ge0/6	0	0	0
ge0/7	0	0	0
ge0/8	0	0	0
ge0/9	0	0	0
ge0/10	0	0	0

Figure 4.183 Switching > DHCP Snooping > Base > Statistics

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row.
MAC Verify Failures	The number of DHCP messages that were dropped because the source MAC address and client hardware address did not match. MAC address verification is performed only if it is globally enabled.
Client Ifc Mismatch	The number of packets that were dropped by DHCP snooping because the interface and VLAN on which the packet was received does not match the client's interface and VLAN information stored in the binding database.
DHCP Server Msgs Received	The number of DHCP server messages (DHCPOFFER, DHCPACK, DHCPNAK, DHCPRELEASEQUERY) that have been dropped on an untrusted port.
Refresh	Click Refresh to update the screen.
Clear Counters	Click Clear Counters to reset all statistics to zero for all interfaces.

4.4.3.2 L2 Relay

Global

Use the DHCP L2 Relay Global Configuration page to control the administrative mode of DHCP Layer 2 Relay on the device. In Layer 2 switched networks, there may be one or more infrastructure devices (for example, a switch) between the client and the L3 Relay agent/DHCP server. In this instance, some of the client device information required by the L3 Relay agent may not be visible to it. When this happens, an L2 Relay agent can be used to add the information that the L3 Relay Agent and DHCP server need to perform their roles in IP address configuration and assignment.

To access this page, click **Switching > DHCP Snooping > L2 Relay > Global**.




Figure 4.184 Switching > DHCP Snooping > L2 Relay > Global

The following table describes the items in the previous figure.

Item	Description
Admin Mode	The global mode of DHCP L2 relay on the device. When enabled, the device can act as a DHCP L2 relay agent. This functionality must also be enabled on each port you want this service to operate on.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

Interface Configuration

Use the DHCP L2 Relay Interface Configuration page to enable L2 DHCP relay on individual ports. Note that L2 DHCP relay must also be enabled globally on the device. To change the DHCP L2 relay settings for one or more interfaces, select each entry to modify and click Edit. The same settings are applied to all selected interfaces.

To access this page, click **Switching > DHCP Snooping > L2 Relay > Interface Configuration**.



Figure 4.185 Switching > DHCP Snooping > L2 Relay > Interface Configuration

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. When configuring the settings for one or more interfaces, this field identifies each interface that is being configured.

Item	Description
L2 Relay Mode	The administrative mode of L2 relay mode on the interface. When enabled, the interface can act as a DHCP relay agent and add information that the L3 relay agent and DHCP server need to perform their roles in IP address configuration and assignment.
Trust Mode	The L2 relay trust mode for the interface, which is one of the following: <ul style="list-style-type: none"> Trusted: A trusted interface usually connects to other agents or servers participating in the DHCP interaction (e.g. other L2 or L3 relay agents or servers). An interface in this mode always expects to receive DHCP packets that include Option 82 information. If Option 82 information is not included, these packets are discarded. Untrusted: An untrusted interface is generally connected to clients. DHCP packets arriving on an untrusted interface are never expected to carry Option 82 and are discarded if they do.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.

VLAN Configuration

Use the DHCP L2 Relay VLAN Configuration page to control the DHCP L2 relay settings on a particular VLAN. The VLAN is identified by a service VLAN ID (S-VID), which a service provider uses to identify a customer's traffic while traversing the provider network to multiple remote sites. The device uses the VLAN membership of the switch port client (the customer VLAN ID, or C-VID) to perform a lookup on a corresponding S-VID.

To access this page, click **Switching > DHCP Snooping > L2 Relay > VLAN Configuration**.



Figure 4.186 Switching > DHCP Snooping > L2 Relay > VLAN Configuration

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The VLAN associated with the rest of the data in the row. When configuring the settings for one or more VLANs, this field identifies each VLAN that is being configured.
Circuit ID	The administrative mode of the circuit ID. When enabled, if a client sends a DHCP request to the device and the client is in a VLAN that corresponds to the S-VID, the device adds the client's interface number to the Circuit ID sub-option of Option 82 in the DHCP request packet. This enables the device to reduce the broadcast domain to which the server replies are switched when the broadcast bit is set for DHCP packets. When this bit is set, the server is required to echo Option 82 in replies. Since the circuit-id field contains the client interface number, the L2 relay agent can forward the response to the requesting interface only, rather to all ports in the VLAN).
Remote ID	The DHCP remote identifier string. When a string is entered here, if a client sends a DHCP request to the device and the client is in a VLAN that corresponds to the S-VID, the device adds the string to the Remote-ID suboption of Option 82 in the DHCP request packet. This suboption can be used by the server for parameter assignment. The content of this option is vendor-specific.

Item	Description
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new DHCP L2 relay VLAN configuration. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

To add a new DHCP L2 relay VLAN configuration:

Click **Switching > DHCP Snooping > L2 Relay > VLAN Configuration > Add**.

Figure 4.187 Switching > DHCP Snooping > L2 Relay > VLAN Configuration > Add

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The VLAN associated with the rest of the data in the row. When configuring the settings for one or more VLANs, this field identifies each VLAN that is being configured.
Circuit ID	The administrative mode of the circuit ID. When enabled, if a client sends a DHCP request to the device and the client is in a VLAN that corresponds to the S-VID, the device adds the client's interface number to the Circuit ID sub-option of Option 82 in the DHCP request packet. This enables the device to reduce the broadcast domain to which the server replies are switched when the broadcast bit is set for DHCP packets. When this bit is set, the server is required to echo Option 82 in replies. Since the circuit-id field contains the client interface number, the L2 relay agent can forward the response to the requesting interface only, rather than to all ports in the VLAN).
Remote ID	The DHCP remote identifier string. When a string is entered here, if a client sends a DHCP request to the device and the client is in a VLAN that corresponds to the S-VID, the device adds the string to the Remote-ID suboption of Option 82 in the DHCP request packet. This suboption can be used by the server for parameter assignment. The content of this option is vendor-specific.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

Statistics

The DHCP L2 Relay Interface Statistics page shows statistical information about the L2 DHCP Relay requests received on trusted and untrusted interfaces. An interface is untrusted when the DHCP L2 relay trust mode is disabled.

To access this page, click **Switching > DHCP Snooping > L2 Relay > Statistics**.

Interface	Untrusted Server Messages With Option 82	Untrusted Client Messages With Option 82	Trusted Server Messages With Option 82	Trusted Client Messages With Option 82
ge0/1	0	0	0	0
ge0/2	0	0	0	0
ge0/3	0	0	0	0
ge0/4	0	0	0	0
ge0/5	0	0	0	0
ge0/6	0	0	0	0
ge0/7	0	0	0	0
ge0/8	0	0	0	0
ge0/9	0	0	0	0
ge0/10	0	0	0	0

Figure 4.188 Switching > DHCP Snooping > L2 Relay > Statistics

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row.
Untrusted Server Messages With Option-82	The number of messages received on an untrusted interface from a DHCP server that contained Option 82 data. These messages are dropped.
Untrusted Client Messages With Option-82	The number of messages received on an untrusted interface from a DHCP client that contained Option 82 data. These messages are dropped.
Trusted Server Messages With Option-82	The number of messages received on a trusted interface from a DHCP server that contained Option 82 data. These messages are forwarded.
Trusted Client Messages With Option-82	The number of messages received on a trusted interface from a DHCP client that contained Option 82 data. These messages are forwarded.
Refresh	Click Refresh to update the screen.
Clear Counters	Click Clear Counters to reset all counters to zero.

4.4.3.3 IP Source Guard

Interface Configuration

Use the Interface Configuration page to configure IP Source Guard (IPSG) on each interface. IPSG is a security feature that filters IP packets based on source ID. This feature helps protect the network from attacks that use IP address spoofing to compromise or overwhelm the network. The source ID may be either the source IP address or a {source IP address, source MAC address} pair. The DHCP snooping bindings database, along with IPSG entries in the database, identify authorized source IDs. If you enable IPSG on a port where DHCP snooping is disabled or where DHCP snooping is enabled but the port is trusted, all IP traffic received on that port is dropped depending on the admin-configured IPSG entries. Additionally, IPSG interacts with port security, also known as port MAC locking, to enforce the source MAC address in received packets. Port security controls source MAC address learning in the Layer 2 forwarding database (MAC address table). When a frame is received with a previously unlearned source MAC address, port security queries the IPSG feature to determine whether the MAC address belongs to a valid binding. To change the IPSG configuration on one or more interfaces, select each entry to modify and click Edit. The same settings are applied to all selected interfaces.

To access this page, click **Switching > DHCP Snooping > IP Source Guard > Interface Configuration**.

Interface	IP Source Guard	Port Security
0/1	Disabled	Disabled
0/2	Disabled	Disabled
0/3	Disabled	Disabled
0/4	Disabled	Disabled
0/5	Disabled	Disabled
0/6	Disabled	Disabled
0/7	Disabled	Disabled
0/8	Disabled	Disabled
0/9	Disabled	Disabled
0/10	Disabled	Disabled

Figure 4.189 Switching > DHCP Snooping > IP Source Guard > Interface Configuration

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. When configuring the settings for one or more interfaces in the Edit DHCP Snooping IP Source Guard Interface Configuration window, this field identifies each interface that is being configured.
IP Source Guard	The administrative mode of IPSG on the interface. When enabled, the source IP address is validated against the DHCP snooping bindings database, and DHCP packets will not be forwarded if the sender's IP address is not in the DHCP snooping bindings database.
Port Security	The administrative mode of IPSG Port Security on the interface. When IPSG Port Security is enabled, the packets will not be forwarded if the sender MAC address is not the in forwarding database table or the DHCP snooping bindings database. To enforce filtering based on MAC address, Port Security must be enabled globally and on the interface. IPSG Port Security cannot be enabled if IPSG is disabled.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.

Bindings

To access this page, click **Switching > DHCP Snooping > IP Source Guard > Bindings**.

Interface	VLAN ID	MAC Address	IP Address	Filter Type	Binding Type
Table is Empty					

Figure 4.190 Switching > DHCP Snooping > IP Source Guard > Bindings

The following table describes the items in the previous figure.

Item	Description
Interface	The interface on which the source ID is authorized.
VLAN ID	The authorized ingress VLAN for the binding rule.
MAC Address	The authorized sender MAC address for the binding rule.
IP Address	The authorized source IP address for the binding rule.

Item	Description
Filter Type	The IPSG filter type, which is one of the following: <ul style="list-style-type: none"> ■ IP: Only the IP address configured in the binding is used as the source ID to filter IP packets. ■ IP-MAC: Both the IP address and its associated MAC address are used to verify whether the IP packets are allowed on the interface. The MAC address is used for IP source enforcement when IPSG port security is enabled on the interface.
Binding Type	The binding type, which is either learned through DHCP snooping (Dynamic) or statically configured in the IPSG bindings database by an administrator (Static).
Refresh	Click Refresh to update the screen.
Add	Click Add to configure a static IPSG entry in the bindings database. See the following procedure.
Remove	Click Remove to remove the selected entries.

To configure a static IPSG entry in the bindings database:

Click **Switching > DHCP Snooping > IP Source Guard > Bindings > Add**.

Figure 4.191 Switching > DHCP Snooping > IP Source Guard > Bindings > Add

The following table describes the items in the previous figure.

Item	Description
Interface	The interface on which the source ID is authorized.
VLAN ID	The authorized ingress VLAN for the binding rule.
MAC Address	The authorized sender MAC address for the binding rule.
IP Address	The authorized source IP address for the binding rule.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.4.4 IPv6 DHCP Snooping

4.4.4.1 Base

Global

Use the IPv6 DHCP Snooping Configuration page to view and configure the global settings for IPv6 DHCP snooping. IPv6 DHCP snooping is a security feature that monitors DHCPv6 messages between a DHCPv6 client and DHCPv6 servers to filter harmful DHCPv6 messages and to build a bindings database of {MAC address, IPv6 address, VLAN ID, port} tuples that are considered authorized. You can enable IPv6 DHCP snooping globally and on specific VLANs, and configure ports within the VLAN to be trusted or untrusted. If a DHCPv6 message arrives on an untrusted port, IPv6 DHCP snooping filters messages that are not from authorized DHCPv6 clients. DHCPv6 server messages are forwarded only through trusted ports.

To access this page, click **Switching > IPv6 DHCP Snooping > Base > Global**.

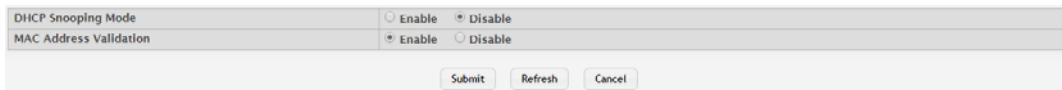


Figure 4.192 Switching > IPv6 DHCP Snooping > Base > Global

The following table describes the items in the previous figure.

Item	Description
DHCP Snooping Mode	The administrative mode of IPv6 DHCP snooping on the device.
MAC Address Validation	Enables or Disables the verification of the sender MAC address for IPv6 DHCP snooping. When enabled, the device checks packets that are received on untrusted interfaces to verify that the MAC address and the DHCPv6 client hardware address match. If the addresses do not match, the device drops the packet.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

VLAN Configuration

Use the IPv6 DHCP Snooping VLAN Configuration page to view and configure the IPv6 DHCP snooping settings on VLANs that exist on the device. IPv6 DHCP snooping can be configured on switching VLANs and routing VLANs. For Layer 2 (non-routing) VLANs, IPv6 DHCP snooping forwards valid DHCPv6 client messages received on the VLANs. The message is forwarded on all trusted interfaces in the VLAN. When a DHCPv6 packet is received on a routing VLAN, the IPv6 DHCP snooping application applies its filtering rules and updates the bindings database. If a client message passes filtering rules, the message is placed into the software forwarding path, where it may be processed by the DHCPv6 relay agent, the local DHCPv6 server, or forwarded as an IPv6 packet.

To access this page, click **Switching > IPv6 DHCP Snooping > Base > VLAN Configuration**.

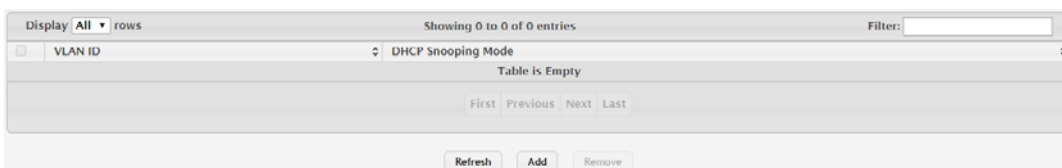


Figure 4.193 Switching > IPv6 DHCP Snooping > Base > VLAN Configuration

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The VLAN ID that is enabled for IPv6 DHCP snooping. In the Add IPv6 DHCP Snooping VLAN Configuration window, this field lists the VLAN ID of all VLANs that exist on the device.
DHCP Snooping Mode	The current administrative mode of IPv6 DHCP snooping for the VLAN. Only VLANs that are enabled for IPv6 DHCP snooping appear in the list.
Refresh	Click Refresh to update the screen.
Add	Click Add to enable a VLAN for IPv6 DHCP snooping. To select multiple VLANs, CTRL + click each VLAN to select. See the following procedure.
Remove	Click Remove to disable IPv6 DHCP snooping for the selected entries.

To enable a VLAN for IPv6 DHCP snooping:

Click **Switching > IPv6 DHCP Snooping > Base > VLAN Configuration > Add**.

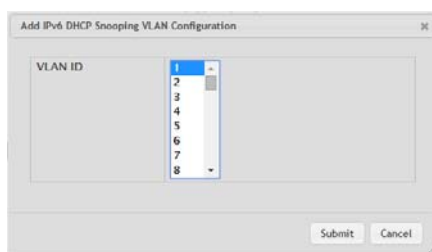


Figure 4.194 Switching > IPv6 DHCP Snooping > Base > VLAN Configuration > Add

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The VLAN ID that is enabled for IPv6 DHCP snooping. In the Add IPv6 DHCP Snooping VLAN Configuration window, this field lists the VLAN ID of all VLANs that exist on the device.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

Interface Configuration

Use the IPv6 DHCP Snooping Interface Configuration page to view and configure the IPv6 DHCP snooping settings for each interface. The IPv6 DHCP snooping feature processes incoming DHCPv6 messages. For RELEASE and DECLINE messages, the feature compares the receive interface and VLAN with the client's interface and VLAN in the binding database. If the interfaces do not match, the application logs the event (when logging of invalid packets is enabled) and drops the message. If MAC address validation is globally enabled, messages that pass the initial validation are checked to verify that the source MAC address and the DHCPv6 client hardware address match. Where there is a mismatch, IPv6 DHCP snooping logs the event (when logging of invalid packets is enabled) and drops the packet.

To access this page, click **Switching > IPv6 DHCP Snooping > Base > Interface Configuration**.

The screenshot shows a web interface for configuring IPv6 DHCP Snooping. At the top, it says 'Display 10 rows' and 'Showing 1 to 10 of 22 entries'. Below this is a table with columns: Interface, Trust State, Log Invalid Packets, Rate Limit (pps), and Burst Interval (Seconds). The table lists interfaces from ge0/1 to ge0/10, all with 'Disabled' Trust State and 'Disabled' Log Invalid Packets. At the bottom of the table are navigation buttons: First, Previous, 1, 2, 3, Next, Last. Below the table are 'Refresh' and 'Edit' buttons.

Interface	Trust State	Log Invalid Packets	Rate Limit (pps)	Burst Interval (Seconds)
ge0/1	Disabled	Disabled		
ge0/2	Disabled	Disabled		
ge0/3	Disabled	Disabled		
ge0/4	Disabled	Disabled		
ge0/5	Disabled	Disabled		
ge0/6	Disabled	Disabled		
ge0/7	Disabled	Disabled		
ge0/8	Disabled	Disabled		
ge0/9	Disabled	Disabled		
ge0/10	Disabled	Disabled		

Figure 4.195 Switching > IPv6 DHCP Snooping > Base > Interface Configuration

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. When configuring the settings for one or more interfaces, this field identifies each interface that is being configured.
Trust State	<p>The trust state configured on the interface. The trust state is one of the following:</p> <ul style="list-style-type: none"> ■ Disabled: The interface is considered to be untrusted and could potentially be used to launch a network attack. DHCPv6 server messages are checked against the bindings database. On untrusted ports, IPv6 DHCP snooping enforces the following security rules: <ul style="list-style-type: none"> – DHCPv6 packets from a DHCPv6 server (ADVERTISE, REPLY, and RECONFIGURE) are dropped. – RELEASE and DECLINE messages are dropped if the MAC address is in the snooping database but the binding's interface is other than the interface where the message was received. – DHCPv6 packets are dropped when the source MAC address does not match the client hardware address if MAC Address Validation is globally enabled. ■ Enabled: The interface is considered to be trusted and forwards DHCPv6 server messages without validation.
Log Invalid Packets	The administrative mode of invalid packet logging on the interface. When enabled, the IPv6 DHCP snooping feature generates a log message when an invalid packet is received and dropped by the interface.

Item	Description
Rate Limit (pps)	The rate limit value for DHCPv6 packets received on the interface. To prevent DHCPv6 packets from being used as a DoS attack when IPv6 DHCP snooping is enabled, the snooping application enforces a rate limit for DHCPv6 packets received on untrusted interfaces. If the incoming rate of DHCPv6 packets exceeds the value of this object during the amount of time specified for the burst interval, the port will be shutdown. You must administratively enable the port to allow it to resume traffic forwarding.
Burst Interval (Seconds)	The burst interval value for rate limiting on this interface. If the rate limit is unspecified, then burst interval has no meaning.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.

Static Bindings

Use the IPv6 DHCP Snooping Static Bindings page to view, add, and remove static bindings in the IPv6 DHCP snooping bindings database.

To access this page, click **Switching > IPv6 DHCP Snooping > Base > Static Bindings**.




Figure 4.196 Switching > IPv6 DHCP Snooping > Base > Static Bindings

The following table describes the items in the previous figure.

Item	Description
Interface	The interface on which the DHCPv6 client is authorized.
MAC Address	The MAC address associated with the DHCP client. This is the key to the binding database.
VLAN ID	The ID of the VLAN the client is authorized to use.
IP Address	The IPv6 address of the client.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new static entry to the IPv6 DHCP snooping bindings table. See the following procedure.
Remove	Click Remove to remove the selected entries.

To add a new static entry to the IPv6 DHCP snooping bindings table:

Click **Switching > IPv6 DHCP Snooping > Base > Static Bindings > Add**.




Figure 4.197 Switching > IPv6 DHCP Snooping > Base > Static Bindings > Add

The following table describes the items in the previous figure.

Item	Description
Interface	The interface on which the DHCPv6 client is authorized.

Item	Description
MAC Address	The MAC address associated with the DHCP client. This is the key to the binding database.
VLAN ID	The ID of the VLAN the client is authorized to use.
IP Address	The IPv6 address of the client.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

Dynamic Bindings

Use the IPv6 DHCP Snooping Dynamic Bindings page to view and clear dynamic bindings in the IPv6 DHCP snooping bindings database. The IPv6 DHCP snooping feature uses DHCPv6 messages to build and maintain the bindings database. The bindings database includes data for clients only on untrusted ports. IPv6 DHCP snooping creates a tentative binding from DHCPv6 SOLICIT and REQUEST messages. Tentative bindings tie a client to an interface (the interface where the DHCPv6 client message was received). Tentative bindings are completed when IPv6 DHCP snooping learns the client's IPv6 address from a REPLY message on a trusted port. DHCP snooping removes bindings in response to DECLINE and RELEASE messages.

To access this page, click **Switching > IPv6 DHCP Snooping > Base > Dynamic Bindings**.




Figure 4.198 Switching > IPv6 DHCP Snooping > Base > Dynamic Bindings

The following table describes the items in the previous figure.

Item	Description
Interface	The interface on which the DHCPv6 client message was received.
MAC Address	The MAC address associated with the DHCPv6 client that sent the message. This is the key to the binding database.
VLAN ID	The VLAN ID of the client interface.
IP Address	The IPv6 address assigned to the client by the DHCPv6 server.
Lease Time	The remaining IPv6 address lease time for the client.
Refresh	Click Refresh to update the screen.
Clear	Click Clear to remove the selected entries in the database.

Persistent

Use the IPv6 DHCP Snooping Persistent Configuration page to configure the persistent location of the IPv6 DHCP snooping bindings database. The bindings database can be stored locally on the device or on a remote system somewhere else in the network. The device must be able to reach the IP address of the remote system to send bindings to a remote database.

To access this page, click **Switching > IPv6 DHCP Snooping > Base > Persistent**.



Figure 4.199 Switching > IPv6 DHCP Snooping > Base > Persistent

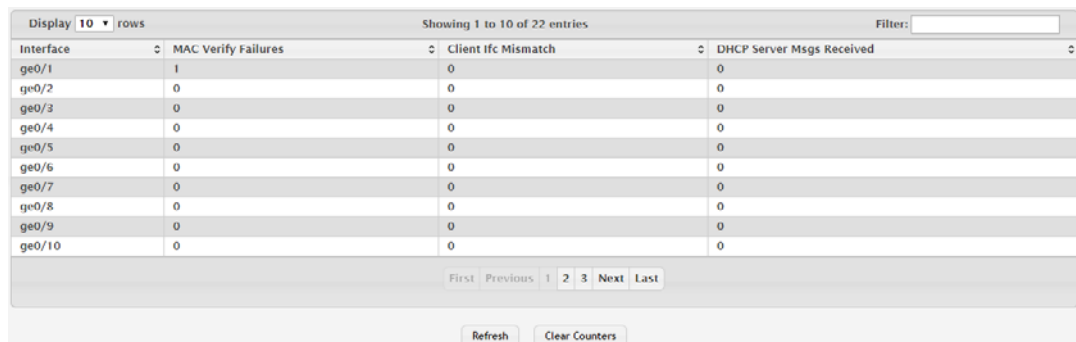
The following table describes the items in the previous figure.

Item	Description
Store	The location of the IPv6 DHCP snooping bindings database, which is either locally on the device (Local) or on a remote system (Remote).
Remote IP Address	The IP address of the system on which the IPv6 DHCP snooping bindings database will be stored. This field is available only if Remote is selected in the Store field.
Remote File Name	The file name of the IPv6 DHCP snooping bindings database in which the bindings are stored. This field is available only if Remote is selected in the Store field.
Write Delay (Seconds)	The amount of time to wait between writing bindings information to persistent storage. This allows the device to collect as many entries as possible (new and removed) before writing them to the persistent file.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

Statistics

Use the IPv6 DHCP Snooping Statistics page to view and clear per-interface statistics about the DHCPv6 messages filtered by the IPv6 DHCP snooping feature. Only interfaces that are enabled for IPv6 DHCP snooping and are untrusted appear in the table.

To access this page, click **Switching > IPv6 DHCP Snooping > Base > Statistics**.



Interface	MAC Verify Failures	Client Ifc Mismatch	DHCP Server Msgs Received
ge0/1	0	0	0
ge0/2	0	0	0
ge0/3	0	0	0
ge0/4	0	0	0
ge0/5	0	0	0
ge0/6	0	0	0
ge0/7	0	0	0
ge0/8	0	0	0
ge0/9	0	0	0
ge0/10	0	0	0

Figure 4.200 Switching > IPv6 DHCP Snooping > Base > Statistics

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row.
MAC Verify Failures	The number of DHCPv6 messages that were dropped because the source MAC address and client hardware address did not match. MAC address verification is performed only if it is globally enabled.
Client Ifc Mismatch	The number of packets that were dropped by IPv6 DHCP snooping because the interface and VLAN on which the packet was received does not match the client's interface and VLAN information stored in the binding database.
DHCP Server Msgs Received	The number of DHCPv6 server messages (ADVERTISE, REPLY, RECONFIGURE, RELAY-REPL) that have been dropped on an untrusted port.
Refresh	Click Refresh to update the screen.
Clear Counters	Click Clear Counters to reset all counters to zero.

4.4.5 DVLAN

DVLAN Tunneling allows the use of a second tag on network traffic. The additional tag helps differentiate between customers in the Metropolitan Area Networks (MAN) while preserving individual customer's VLAN identification when they enter their own 802.1Q domain.

With the introduction of this second tag, you do not need to divide the 4k VLAN ID space to send traffic on an Ethernet-based MAN.

With DVLAN Tunneling enabled, every frame that is transmitted from an interface has a new VLAN tag (S-tag) attached while every packet that is received from an interface has a VLAN tag (S-tag) removed (if one or more tags are present).

DVLAN also supports up to 4 Tag Protocol Identifier (TPID) values per switch and the ability to map these values to ports. This allows you to configure the same or different TPIDs for different ports.

4.4.5.1 Configuration

The DVLAN Configuration page allows you to configure the Tag Protocol Identifier (TPID) to include in frames transmitted by interfaces that are enabled for double VLAN (DVLAN) tagging. DVLAN tagging allows the device to add a second (outer) VLAN tag to the frame while preserving the original (inner) VLAN tagging information.

To access this page, click **Switching > DVLAN > Configuration**.





Figure 4.201 Switching > DVLAN > Configuration

The following table describes the items in the previous figure.

Item	Description
Primary TPID	The two-byte hex EtherType value to be used as the first 16 bits of the DVLAN tag. The value configured in this field is used as the primary TPID for all interfaces that are enabled for DVLAN tagging. The Primary TPID can be one of the following: <ul style="list-style-type: none">■ 0x8100: IEEE 802.1Q customer VLAN tag type■ 0x88a8: Virtual Metropolitan Area Network (VMAN) tag type■ Custom Tag: User-defined EtherType value
Secondary TPIDs	The two-byte hex EtherType values available to be configured as secondary TPIDs. Only the options you configure as Secondary TPIDs can be selected as the Primary TPID. To add Secondary TPIDs to the list, click  button and select one or more of the following options: <ul style="list-style-type: none">■ 802.1Q Tag: IEEE 802.1Q customer VLAN tag type, represented by the EtherType value 0x8100. This value indicates that the frame includes a VLAN tag. If this value is already configured as a primary or secondary TPID, it cannot be selected.■ vMAN Tag: Virtual Metropolitan Area Network (VMAN) tag type, represented by the EtherType value 0x88a8. This value indicates that the frame is DVLAN tagged. If this value is already configured as a primary or secondary TPID, it cannot be selected.■ Custom Tag: User-defined EtherType value. If you select this option, specify the EtherType value in the available field.
Refresh	Click Refresh to update the screen.

4.4.5.2 Summary

The DVLAN Summary page allows you to view the Global and Default TPIDs configured for all ports on the system.

To access this page, click **Switching > DVLAN > Summary**.

Primary TPID	0x8100
Secondary TPIDs	
Refresh	

Figure 4.202 Switching > DVLAN > Summary

The following table describes the items in the previous figure.

Item	Description
Primary TPID	The two-byte hex EtherType value used as the first 16 bits of the DVLAN tag. This value identifies the frame as one of the following types: <ul style="list-style-type: none"> 0x8100: IEEE 802.1Q VLAN tag type. This value indicates that the frame includes a VLAN tag. 0x88a8: Virtual Metropolitan Area Network (VMAN) tag type. This value indicates that the frame is double VLAN tagged. Custom Tag: Any TPID value other than 0x8100 or 0x88a8 is a user-defined EtherType value.
Secondary TPIDs	The two-byte hex EtherType values configured as secondary TPIDs.
Refresh	Click Refresh to update the screen.

4.4.5.3 Interface Summary

Use the DVLAN Interface Summary page to view and configure the double VLAN (DVLAN) tag settings for each interface. Double VLAN tagging allows service providers to create Virtual Metropolitan Area Networks (VMANs). With DVLAN tagging, service providers can pass VLAN traffic from one customer domain to another through a metro core. By using an additional tag on the traffic, the interface can differentiate between customers in the MAN while preserving an individual customer's VLAN identification that is used when the traffic enters the customer's 802.1Q domain.

To access this page, click **Switching > DVLAN > Interface Summary**.

Interface	Interface Mode	Interface EtherType
ge0/1	Disable	0x8100
ge0/2	Disable	0x8100
ge0/3	Disable	0x8100
ge0/4	Disable	0x8100
ge0/5	Disable	0x8100
ge0/6	Disable	0x8100
ge0/7	Disable	0x8100
ge0/8	Disable	0x8100
ge0/9	Disable	0x8100
ge0/10	Disable	0x8100

Figure 4.203 Switching > DVLAN > Interface Summary

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row.

Item	Description
Interface Mode	The administrative mode of double VLAN tagging on the interface. When DVLAN tagging is enabled, every frame that is transmitted from the interface has a DVLAN tag attached, and every packet that is received from the interface has a tag removed (if one or more tags are present).
Interface EtherType	The EtherType value to be used as the first 16 bits of the DVLAN tag. If one or more secondary TPIDs have been configured for the interface, these EtherType values are also displayed.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.

4.4.6 Dynamic ARP Inspection

Dynamic ARP Inspection (DAI) is a security feature that rejects invalid and malicious ARP packets. DAI prevents a class of man-in-the-middle attacks, where an unfriendly station intercepts traffic for other stations by poisoning the ARP caches of its unsuspecting neighbors. The miscreant sends ARP requests or responses mapping another station's IP address to its own MAC address.

DAI relies on DHCP snooping. DHCP snooping listens to DHCP message exchanges and builds a binding database of valid {MAC address, IP address, VLAN, and interface} tuples.

When DAI is enabled, the switch drops ARP packets whose sender MAC address and sender IP address do not match an entry in the DHCP snooping bindings database. You can optionally configure additional ARP packet validation.

4.4.6.1 Global

Use the Global Configuration page to configure global DAI settings.

To access this page, click **Switching > Dynamic ARP Inspection > Global**.

Figure 4.204 Switching > Dynamic ARP Inspection > Global

The following table describes the items in the previous figure.

Item	Description
Validate Source MAC	When this option is selected, DAI verifies that the sender hardware address in the ARP packet equals the source MAC address in the Ethernet header. If the addresses do not match, the ARP packet is dropped.
Validate Destination MAC	When this option is selected, DAI verifies that the target hardware address in the ARP packet equals the destination MAC address in the Ethernet header. If the addresses do not match, the ARP packet is dropped. This check applies only to ARP responses because the target MAC address is unspecified in ARP requests.
Validate IP	When this option is selected, DAI drops ARP packets with an invalid IP address. The following IP addresses are considered invalid: <ul style="list-style-type: none"> ■ 0.0.0.0 ■ 255.255.255.255 ■ All IP multicast addresses ■ All class E addresses (240.0.0.0/4) ■ Loopback addresses (in the range 127.0.0.0/8)

Item	Description
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.4.6.2 VLAN

Use the Dynamic ARP Inspection VLAN Configuration page to view and configure Dynamic ARP Inspection (DAI) settings for VLANs. When DAI is enabled on a VLAN, DAI is enabled on all interfaces that are members of that VLAN.

To access this page, click **Switching > Dynamic ARP Inspection > VLAN**.




Figure 4.205 Switching > Dynamic ARP Inspection > VLAN

The following table describes the items in the previous figure.

Item	Description
VLAN ID	Lists each VLAN that has been enabled for DAI. After you click Add, use the VLAN ID menu to select the VLAN on which to enable DAI. A VLAN does not need to exist on the system to be enabled for DAI.
Log Invalid Packets	Indicates whether DAI logging is enabled on this VLAN. When logging is enabled, DAI generates a log message whenever an invalid ARP packet is discovered and dropped.
ARP ACL Name	The name of the of ARP access control list (ACL) that the VLAN uses as the filter for ARP packet validation. The ARP ACL must already exist on the system to associate it with a DAI-enabled VLAN. ARP ACLs include permit rules only.
Static	Determines whether to use the DHCP snooping database for ARP packet validation if the packet does not match any ARP ACL rules. The options are as follows: <ul style="list-style-type: none"> ■ Enable: The ARP packet will be validated by the ARP ACL rules only. Packets that do not match any ARP ACL rules are dropped without consulting the DHCP snooping database. ■ Disable: The ARP packet needs further validation by using the entries in the DHCP Snooping database.
Refresh	Click Refresh to update the screen.
Add	Click Add to enable DAI on a VLAN and configure the optional DAI settings. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to disable DAI for the selected entries.

To DAI on a VLAN and configure the optional DAI settings:
Click **Switching > Dynamic ARP Inspection > VLAN > Add**.

Figure 4.206 Switching > Dynamic ARP Inspection > VLAN > Add

The following table describes the items in the previous figure.

Item	Description
VLAN ID	Lists each VLAN that has been enabled for DAI. After you click Add, use the VLAN ID menu to select the VLAN on which to enable DAI. A VLAN does not need to exist on the system to be enabled for DAI.
Log Invalid Packets	Indicates whether DAI logging is enabled on this VLAN. When logging is enabled, DAI generates a log message whenever an invalid ARP packet is discovered and dropped.
ARP ACL Name	The name of the of ARP access control list (ACL) that the VLAN uses as the filter for ARP packet validation. The ARP ACL must already exist on the system to associate it with a DAI-enabled VLAN. ARP ACLs include permit rules only.
Static	Determines whether to use the DHCP snooping database for ARP packet validation if the packet does not match any ARP ACL rules. The options are as follows: <ul style="list-style-type: none"> ■ Enable: The ARP packet will be validated by the ARP ACL rules only. Packets that do not match any ARP ACL rules are dropped without consulting the DHCP snooping database. ■ Disable: The ARP packet needs further validation by using the entries in the DHCP Snooping database.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.4.6.3 Interface

Use the Interface Configuration page to configure the per-interface Dynamic ARP Inspection (DAI) settings.

To access this page, click **Switching > Dynamic ARP Inspection > Interface**.

Interface	Trust State	Rate Limit	Burst Interval
ge0/1	Disabled	15	1
ge0/2	Disabled	15	1
ge0/3	Disabled	15	1
ge0/4	Disabled	15	1
ge0/5	Disabled	15	1
ge0/6	Disabled	15	1
ge0/7	Disabled	15	1
ge0/8	Disabled	15	1
ge0/9	Disabled	15	1
ge0/10	Disabled	15	1

Figure 4.207 Switching > Dynamic ARP Inspection > Interface

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. In the Edit Interface Configuration window, this field identifies the interface that is being configured.
Trust State	Indicates whether the DAI feature should check traffic on the interface for possible ARP packet violations. Trust state can be enabled or disabled after you select an interface and click Edit . This field has one of the following values: <ul style="list-style-type: none"> ■ Enabled: The interface is trusted. ARP packets arriving on this interface are forwarded without DAI validation. ■ Disabled: The interface is not trusted. ARP packets arriving on this interface are subjected to ARP inspection.
Rate Limit	The maximum rate for incoming ARP packets on the interface, in packets per second (pps). If the incoming rate exceeds the configured limit, the ARP packets are dropped. Rate limiting can be enabled or disabled after you select an interface and click Edit .
Burst Interval	The number of consecutive seconds the interface is monitored for incoming ARP packet rate limit violations.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.

4.4.6.4 ACL Summary

Use the ACL Summary page to configure ARP Access Control Lists (ACLs). An ARP ACL can contain one or more permit rules. Each rule contains the IP address and MAC address of a system allowed to send ARP packets. When an ARP ACL is associated with a DAI-enabled VLAN, and an ARP packet is received on an interface that is a member of that VLAN, DAI validates the address information in the ARP packet against the rules in the ACL. If the sender information in the ARP packet matches a rule in the ARP ACL, DAI considers the packet to be valid, and the packet is forwarded.

To access this page, click **Switching > Dynamic ARP Inspection > ACL Summary**.




Figure 4.208 Switching > Dynamic ARP Inspection > ACL Summary

The following table describes the items in the previous figure.

Item	Description
ACL Name	The name of the ACL. Only the ACLs that appear in this column can be referenced by DNI-enabled VLANs.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new ARP ACL. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

To add a new ARP ACL:

Click **Switching > Dynamic ARP Inspection > ACL Summary > Add**.



Figure 4.209 Switching > Dynamic ARP Inspection > ACL Summary > Add

The following table describes the items in the previous figure.

Item	Description
ACL Name	The name of the ACL. Only the ACLs that appear in this column can be referenced by DNI-enabled VLANs.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.4.6.5 ACL Configuration

Use the ACL Configuration page to configure rules for the existing ARP Access Control Lists (ACLs) on the system and to view summary information about the rules that have been added to an ACL. Each rule contains the IP address and MAC address of a system allowed to send ARP packets.

To access this page, click **Switching > Dynamic ARP Inspection > ACL Configuration**.

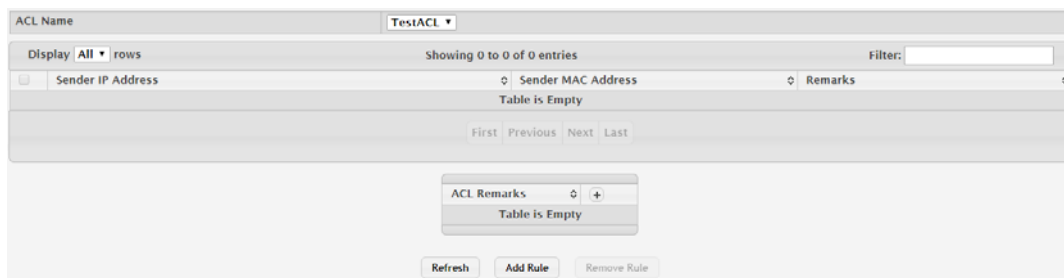


Figure 4.210 Switching > Dynamic ARP Inspection > ACL Configuration

The following table describes the items in the previous figure.

Item	Description
ACL Name	The menu contains the ARP ACL names that exist on the system.
Sender IP Address	The IP address of a system that is permitted to send ARP packets. The ARP packet must match on both the Sender IP Address and Sender MAC Address values in the rule to be considered valid.
Sender MAC Address	The MAC address of a system that is permitted to send ARP packets. The ARP packet must match on both the Sender IP Address and Sender MAC Address values in the rule to be considered valid.
Remarks	One or more remarks configured for the selected ACL and associated with the rule during rule creation.
ACL Remarks	Lists the configured remarks for an ARP ACL. All remarks present in this table are applied to the next rule created with the Add Rule button.
Refresh	Click Refresh to update the screen.
Add Rule	Click Add Rule to add a new rule to an existing ACL. See the following procedure.
Remove Rule	Click Remove Rule to remove the selected entries.

To add a new rule to an existing ACL:

Click **Switching > Dynamic ARP Inspection > ACL Configuration > Add Rule**.

Figure 4.211 Switching > Dynamic ARP Inspection > ACL Configuration > Add Rule

The following table describes the items in the previous figure.

Item	Description
Sender IP Address	The IP address of a system that is permitted to send ARP packets. The ARP packet must match on both the Sender IP Address and Sender MAC Address values in the rule to be considered valid.
Sender MAC Address	The MAC address of a system that is permitted to send ARP packets. The ARP packet must match on both the Sender IP Address and Sender MAC Address values in the rule to be considered valid.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.4.6.6 Statistics

Use the Dynamic ARP Inspection Statistics page to view information about the number of ARP packets that have been forwarded or dropped after being processed by the Dynamic ARP Inspection (DAI) feature. The statistics are shown for each DAI-enabled VLAN.

To access this page, click **Switching > Dynamic ARP Inspection > Statistics**.

Figure 4.212 Switching > Dynamic ARP Inspection > Statistics

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The DAI-enabled VLAN associated with the rest of the information in the row. When DAI is enabled on a VLAN, DAI is enabled on all interfaces that are members of that VLAN.
DHCP Drops	The number of ARP packets that have been dropped by DAI because no matching DHCP snooping binding entry was found in the DHCP snooping database.
ACL Drops	The number of ARP packets that have been dropped by DAI because the sender IP address and sender MAC address in the ARP packet did not match any rules in the ARP ACL associated with this VLAN. The static flag on this VLAN is enabled, which means ARP packets that fail to match an ARP ACL rule are dropped immediately and are not checked against the DHCP snooping database for further validation.
DHCP Permits	The number of ARP packets that were forwarded by DAI because a matching DHCP snooping binding entry was found in the DHCP snooping database.

Item	Description
ACL Permits	The number of ARP packets that were forwarded by DAI because the sender IP address and sender MAC address in the ARP packet matched a rule in the ARP ACL associated with this VLAN.
Bad Source MAC	The number of ARP packets that were dropped by DAI because the sender MAC address in ARP packet did not match the source MAC address in the Ethernet header.
Bad Dest MAC	The number of ARP packets that were dropped by DAI because the target MAC address in the ARP reply packet did not match the destination MAC address in the Ethernet header.
Invalid IP	The number of ARP packets that were dropped by DAI because the sender IP address in the ARP packet or target IP address in the ARP reply packet was invalid. The following IP addresses are considered invalid: <ul style="list-style-type: none"> ■ 0.0.0.0 ■ 255.255.255.255 ■ All IP multicast addresses ■ All class E addresses (240.0.0.0/4) ■ Loopback addresses (in the range 127.0.0.0/8)
Forwarded	The total number of valid ARP packets forwarded by DAI.
Dropped	The total number of invalid ARP packets dropped by DAI.
Refresh	Click Refresh to update the screen.

4.4.7 Filters

Static MAC filtering allows you to associate a MAC address with a VLAN and set of source ports and destination ports. (The availability of source and destination port filters is subject to platform restrictions). Any packet with a static MAC address in a specific VLAN is admitted only if the ingress port is included in the set of source ports; otherwise the packet is dropped. If admitted, the packet is forwarded to all the ports in the destination list.

4.4.7.1 MAC Filters

Use the Static MAC Filter Summary page to view, create, edit, and remove static MAC filters on the device. A MAC filter is a security mechanism that allows Ethernet frames that match the filter criteria (destination MAC address and VLAN ID) to be received and transmitted only on certain ports.

To access this page, click **Switching > Filters > MAC Filters**.




Figure 4.213 Switching > Filters > MAC Filters

The following table describes the items in the previous figure.

Item	Description
MAC Address	The MAC address of the filter. The destination MAC address of an Ethernet frame must match this value to be considered for the filter. When adding or editing a filter, note that you cannot configure the following MAC addresses in this field: <ul style="list-style-type: none"> ■ 00:00:00:00:00:00 ■ 01:80:C2:00:00:00 to 01:80:C2:00:00:10 ■ 01:80:C2:00:00:20 to 01:80:C2:00:00:2F ■ FF:FF:FF:FF:FF:FF
VLAN ID	The VLAN ID associated with the filter. The VLAN ID is used with the MAC address to fully identify the frames to filter.
Source Members	The port(s) included in the inbound filter. If a frame with the MAC address and VLAN ID combination specified in the filter is received on a port in the Source Members list, it is forwarded to a port in the Destination Members list. If the frame that meets the filter criteria is received on a port that is not in the Source Members list, it is dropped. To add source ports to the filter, select one or more ports from the Available Port List field (CTRL + click to select multiple ports). Then, use the appropriate arrow icon to move the selected ports to the Source Members field.
Destination Members	The port(s) included in the outbound filter. A frame with the MAC address and VLAN ID combination specified in the filter is transmitted only out of ports in the list. To add destination ports to the filter, select one or more ports from the Available Port List field (CTRL + click to select multiple ports). Then, use the appropriate arrow icon to add the selected ports to the Source Members field.
Refresh	Click Refresh to update the screen.
Add	Click Add to enable DAI on a VLAN and configure the optional DAI settings. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to disable DAI for the selected entries.

To enable DAI on a VLAN and configure the optional DAI settings:

Click **Switching > Filters > MAC Filters > Add**.

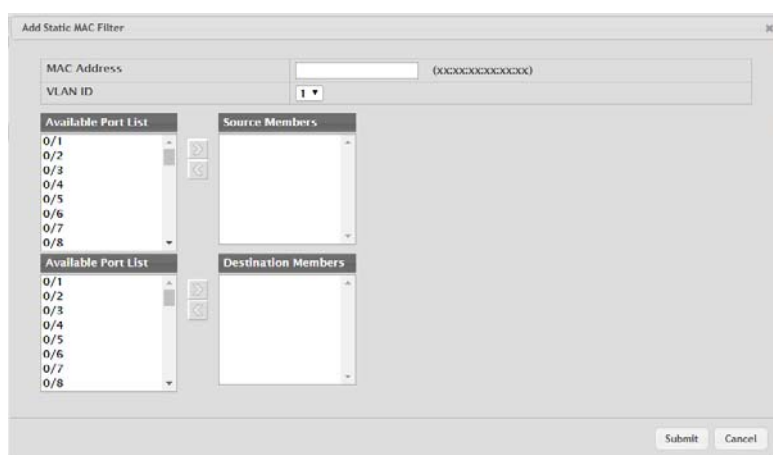


Figure 4.214 Switching > Filters > MAC Filters > Add

The following table describes the items in the previous figure.

Item	Description
MAC Address	The MAC address of the filter. The destination MAC address of an Ethernet frame must match this value to be considered for the filter. When adding or editing a filter, note that you cannot configure the following MAC addresses in this field: <ul style="list-style-type: none"> ■ 00:00:00:00:00:00 ■ 01:80:C2:00:00:00 to 01:80:C2:00:00:10 ■ 01:80:C2:00:00:20 to 01:80:C2:00:00:2F ■ FF:FF:FF:FF:FF:FF
VLAN ID	The VLAN ID associated with the filter. The VLAN ID is used with the MAC address to fully identify the frames to filter.
Source Members	The port(s) included in the inbound filter. If a frame with the MAC address and VLAN ID combination specified in the filter is received on a port in the Source Members list, it is forwarded to a port in the Destination Members list. If the frame that meets the filter criteria is received on a port that is not in the Source Members list, it is dropped. To add source ports to the filter, select one or more ports from the Available Port List field (CTRL + click to select multiple ports). Then, use the appropriate arrow icon to move the selected ports to the Source Members field.
Destination Members	The port(s) included in the outbound filter. A frame with the MAC address and VLAN ID combination specified in the filter is transmitted only out of ports in the list. To add destination ports to the filter, select one or more ports from the Available Port List field (CTRL + click to select multiple ports). Then, use the appropriate arrow icon to add the selected ports to the Source Members field.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.4.8 GARP

4.4.8.1 Switch

Use the GARP Switch Configuration page to set the administrative mode for the features that use the Generic Attribute Registration Protocol (GARP), including GARP VLAN Registration Protocol (GVRP) and GARP Multicast Registration Protocol (GMRP). GARP is a general-purpose protocol that registers any network connectivity or membership-style information. GARP defines a set of switches interested in a given network attribute, such as VLAN ID or multicast address.

To access this page, click **Switching > GARP > Switch**.




Figure 4.215 Switching > GARP > Switch

The following table describes the items in the previous figure.

Item	Description
GVRP Mode	The administrative mode of GVRP on the system. When enabled, GVRP can help dynamically manage VLAN memberships on trunk ports.

Item	Description
GMRP Mode	The administrative mode of GMRP on the system. When enabled, GMRP can help control the flooding of multicast traffic by keeping track of group membership information. GMRP is similar to IGMP snooping in its purpose, but IGMP snooping is more widely used. GMRP must be running on both the host and the switch to function properly.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.4.8.2 Port

Use the GARP Port Configuration page to set the per-interface administrative mode for GARP VLAN Registration Protocol (GVRP) and GARP Multicast Registration Protocol (GMRP). On this page, you can also set the GARP timers for each interface. GVRP and GMRP use the same set of GARP timers to specify the amount of time to wait before transmitting various GARP messages.

To access this page, click **Switching > GARP > Port**.

The screenshot shows a web interface for configuring GARP ports. At the top, it says 'Display 10 rows' and 'Showing 1 to 10 of 22 entries'. There is a search filter box. The table has columns for Interface, GVRP Mode, GMRP Mode, Join Timer (Centiseecs), Leave Timer (Centiseecs), and Leave All Timer (Centiseecs). The rows show interfaces ge0/1 through ge0/10, all with GVRP and GMRP modes set to 'Disabled' and timers set to 20, 60, and 1000 respectively. At the bottom, there are navigation buttons: 'First', 'Previous', '1', '2', '3', 'Next', 'Last', and 'Refresh', 'Edit' buttons.

Interface	GVRP Mode	GMRP Mode	Join Timer (Centiseecs)	Leave Timer (Centiseecs)	Leave All Timer (Centiseecs)
ge0/1	Disabled	Disabled	20	60	1000
ge0/2	Disabled	Disabled	20	60	1000
ge0/3	Disabled	Disabled	20	60	1000
ge0/4	Disabled	Disabled	20	60	1000
ge0/5	Disabled	Disabled	20	60	1000
ge0/6	Disabled	Disabled	20	60	1000
ge0/7	Disabled	Disabled	20	60	1000
ge0/8	Disabled	Disabled	20	60	1000
ge0/9	Disabled	Disabled	20	60	1000
ge0/10	Disabled	Disabled	20	60	1000

Figure 4.216 Switching > GARP > Port

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. When configuring one or more interfaces in the Edit GARP Port Configuration window, this field identifies the interfaces that are being configured.
GVRP Mode	The administrative mode of GVRP on the interface. When enabled, GVRP can help dynamically manage VLAN memberships on trunk ports. GVRP must also be enabled globally for the protocol to be active on the interface. When disabled, the protocol will not be active on the interface, and the GARP timers have no effect.
GMRP Mode	The administrative mode of GMRP on the interface. When enabled, GMRP can help control the flooding of multicast traffic by keeping track of group membership information. GMRP must also be enabled globally for the protocol to be active on the interface. When disabled, the protocol will not be active on the interface, and the GARP timers have no effect.
Join Timer (Centiseecs)	The amount of time between the transmission of GARP PDUs registering (or re-registering) membership for a VLAN or multicast group.

Item	Description
Leave Timer (Centisecs)	The amount of time to wait after receiving an unregister request for a VLAN or multicast group before deleting the associated entry. This timer allows time for another station to assert registration for the same attribute in order to maintain uninterrupted service.
Leave All Timer (Centisecs)	The amount of time to wait before sending a LeaveAll PDU after the GARP application has been enabled on the interface or the last LeaveAll PDU was sent. A LeaveAll PDU indicates that all registrations will shortly be deregistered. Participants will need to rejoin in order to maintain registration.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.

4.4.9 IGMP Snooping

Internet Group Management Protocol (IGMP) Snooping is a feature that allows a switch to forward multicast traffic intelligently on the switch. Multicast IP traffic is traffic that is destined to a host group. Host groups are identified by class D IP addresses, which range from 224.0.0.0 to 239.255.255.255. Based on the IGMP query and report messages, the switch forwards traffic only to the ports that request the multicast traffic. This prevents the switch from broadcasting the traffic to all ports and possibly affecting network performance.

A traditional Ethernet network may be separated into different network segments to prevent placing too many devices onto the same shared media. Bridges and switches connect these segments. When a packet with a broadcast or multicast destination address is received, the switch will forward a copy into each of the remaining network segments in accordance with the IEEE MAC Bridge standard. Eventually, the packet is made accessible to all nodes connected to the network.

This approach works well for broadcast packets that are intended to be seen or processed by all connected nodes. In the case of multicast packets, however, this approach could lead to less efficient use of network bandwidth, particularly when the packet is intended for only a small number of nodes. Packets will be flooded into network segments where no node has any interest in receiving the packet. While nodes will rarely incur any processing overhead to filter packets addressed to un-requested group addresses, they are unable to transmit new packets onto the shared media for the period of time that the multicast packet is flooded. The problem of wasting bandwidth is even worse when the LAN segment is not shared, for example in Full Duplex links.

Allowing switches to snoop IGMP packets is a creative effort to solve this problem. The switch uses the information in the IGMP packets as they are being forwarded throughout the network to determine which segments should receive packets directed to the group address.

4.4.9.1 Configuration

Use the IGMP Snooping Global Configuration and Status page to enable IGMP snooping on the switch and view information about the current IGMP configuration.

To access this page, click **Switching > IGMP Snooping > Configuration**.

The screenshot shows a configuration page for IGMP Snooping. At the top, there is a section for 'Admin Mode' with radio buttons for 'Disable' and 'Enable'. Below this, there are three rows of configuration fields: 'Multicast Control Frame Count' with a value of '0', 'Interfaces Enabled for IGMP Snooping', and 'VLANs Enabled for IGMP Snooping'. At the bottom of the form, there are three buttons: 'Submit', 'Refresh', and 'Cancel'.

Figure 4.217 Switching > IGMP Snooping > Configuration

The following table describes the items in the previous figure.

Item	Description
Admin Mode	The administrative mode of IGMP snooping on the device.
Multicast Control Frame Count	The number of multicast control frames that have been processed by the CPU.
Interfaces Enabled for IGMP Snooping	One or more interfaces on which IGMP snooping is administratively enabled. IGMP snooping must be enabled globally and on an interface for the interface to be able to snoop IGMP packets to determine which segments should receive multicast packets directed to the group address.
VLANs Enabled for IGMP Snooping	One or more VLANs on which IGMP snooping is administratively enabled.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.4.9.2 Interface Configuration

Use the IGMP Snooping Interface Configuration page to configure IGMP snooping settings on specific interfaces. To configure the settings for one or more interfaces, select each entry to modify and click Edit. The same IGMP snooping settings are applied to all selected interfaces.

To access this page, click **Switching > IGMP Snooping > Interface Configuration**.

The screenshot shows a web interface for configuring IGMP snooping on interfaces. At the top, it says 'Display 10 rows' and 'Showing 1 to 10 of 22 entries'. Below this is a table with columns: Interface, Admin Mode, Group Membership Interval, Max Response Time, Multicast Router Expiration Time, and Fast Leave Admin Mode. The table lists interfaces from ge0/1 to ge0/10, all with Admin Mode set to 'Disable', Group Membership Interval set to 260, Max Response Time set to 10, and Multicast Router Expiration Time set to 0. At the bottom of the table, there are navigation buttons: 'First', 'Previous', '2', '3', 'Next', and 'Last'. Below the table are 'Refresh' and 'Edit' buttons.

Interface	Admin Mode	Group Membership Interval	Max Response Time	Multicast Router Expiration Time	Fast Leave Admin Mode
ge0/1	Disable	260	10	0	Disable
ge0/2	Disable	260	10	0	Disable
ge0/3	Disable	260	10	0	Disable
ge0/4	Disable	260	10	0	Disable
ge0/5	Disable	260	10	0	Disable
ge0/6	Disable	260	10	0	Disable
ge0/7	Disable	260	10	0	Disable
ge0/8	Disable	260	10	0	Disable
ge0/9	Disable	260	10	0	Disable
ge0/10	Disable	260	10	0	Disable

Figure 4.218 Switching > IGMP Snooping > Interface Configuration

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. When configuring IGMP snooping settings, this field identifies the interface(s) that are being configured.
Admin Mode	The administrative mode of IGMP snooping on the interface. IGMP snooping must be enabled globally and on an interface for the interface to be able to snoop IGMP packets to determine which segments should receive multicast packets directed to the group address.
Group Membership Interval	The number of seconds the interface should wait for a report for a particular group on the interface before the IGMP snooping feature deletes the interface from the group.
Max Response Time	The number of seconds the interface should wait after sending a query if it does not receive a report for a particular group. The specified value should be less than the Group Membership Interval.

Item	Description
Multicast Router Expiration Time	The number of seconds the interface should wait to receive a query before it is removed from the list of interfaces with multicast routers attached.
Fast Leave Admin Mode	The administrative mode of Fast Leave on the interface. If Fast Leave is enabled, the interface can be immediately removed from the layer 2 forwarding table entry upon receiving an IGMP leave message for a multicast group without first sending out MAC-based general queries.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.

4.4.9.3 Source Specific Multicast

The Source Specific Multicast page displays information about IGMP snooping for source specific multicast.

To access this page, click **Switching > IGMP Snooping > Source Specific Multicast**.




Figure 4.219 Switching > IGMP Snooping > Source Specific Multicast

The following table describes the items in the previous figure.

Item	Description
VLAN ID	VLAN on which the IGMP v3 report is received.
Group	The IPv4 multicast group address.
Interface	The interface on which the IGMP v3 report is received.
Reporter	The IPv4 address of the host that sent the IGMPv3 report.
Source Filter Mode	The source filter mode(Include/Exclude) for the specified group.
Source Address List	List of source IP addresses for which source filtering is requested.
Refresh	Click Refresh to update the screen.

4.4.9.4 VLAN Status

Use the IGMP Snooping VLAN Status page to enable or disable IGMP snooping on system VLANs and to view and configure per-VLAN IGMP snooping settings. Only VLANs that are enabled for IGMP snooping appear in the table.

To access this page, click **Switching > IGMP Snooping > VLAN Status**.




Figure 4.220 Switching > IGMP Snooping > VLAN Status

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The VLAN associated with the rest of the data in the row. When enabling IGMP snooping on a VLAN, use this menu to select the desired VLAN. Only VLANs that have been configured on the system and are not already enabled for IGMP snooping appear in the menu. When modifying IGMP snooping settings, this field identifies the VLAN that is being configured.
Admin Mode	The administrative mode of IGMP snooping on the VLAN. IGMP snooping must be enabled globally and on an VLAN for the VLAN to be able to snoop IGMP packets to determine which network segments should receive multicast packets directed to the group address.
Fast Leave Admin Mode	The administrative mode of Fast Leave on the VLAN. If Fast Leave is enabled, the VLAN can be immediately removed from the layer 2 forwarding table entry upon receiving an IGMP leave message for a multicast group without first sending out MAC-based general queries.
Group Membership Interval (Seconds)	The number of seconds the VLAN should wait for a report for a particular group on the VLAN before the IGMP snooping feature deletes the VLAN from the group.
Max Response Time (Seconds)	The number of seconds the VLAN should wait after sending a query if it does not receive a report for a particular group. The specified value should be less than the Group Membership Interval.
Multicast Router Expiration Time (Seconds)	The number of seconds the VLAN should wait to receive a query before it is removed from the list of VLANs with multicast routers attached.
Refresh	Click Refresh to update the screen.
Add	Click Add to enable IGMP snooping on a VLAN. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to disable IGMP snooping for the selected entries.

To enable IGMP snooping on a VLAN:

Click **Switching > IGMP Snooping > VLAN Status > Add**.

Figure 4.221 Switching > IGMP Snooping > VLAN Status > Add

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The VLAN associated with the rest of the data in the row. When enabling IGMP snooping on a VLAN, use this menu to select the desired VLAN. Only VLANs that have been configured on the system and are not already enabled for IGMP snooping appear in the menu. When modifying IGMP snooping settings, this field identifies the VLAN that is being configured.

Item	Description
Fast Leave Admin Mode	The administrative mode of Fast Leave on the VLAN. If Fast Leave is enabled, the VLAN can be immediately removed from the layer 2 forwarding table entry upon receiving an IGMP leave message for a multicast group without first sending out MAC-based general queries.
Group Membership Interval (Seconds)	The number of seconds the VLAN should wait for a report for a particular group on the VLAN before the IGMP snooping feature deletes the VLAN from the group.
Max Response Time (Seconds)	The number of seconds the VLAN should wait after sending a query if it does not receive a report for a particular group. The specified value should be less than the Group Membership Interval.
Multicast Router Expiration Time (Seconds)	The number of seconds the VLAN should wait to receive a query before it is removed from the list of VLANs with multicast routers attached.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.4.9.5 Multicast Router Configuration

If a multicast router is attached to the switch, its existence can be learned dynamically. You can also statically configure a switch port as a multicast router interface. Use the IGMP Snooping Multicast Router Configuration page to manually configure an interface as a static multicast router interface.

To access this page, click **Switching > IGMP Snooping > Multicast Router Configuration**.

The screenshot shows a web interface for configuring multicast routers. At the top, it says 'Display 10 rows' and 'Showing 1 to 10 of 42 entries'. Below this is a table with two columns: 'Interface' and 'Multicast Router'. The 'Interface' column lists ports from 0/1 to 0/10. The 'Multicast Router' column shows 'Disabled' for all listed interfaces. At the bottom of the table, there are navigation buttons: 'First', 'Previous', '1', '2', '3', '4', '5', 'Next', and 'Last'. Below the table are two buttons: 'Refresh' and 'Edit'.

Interface	Multicast Router
0/1	Disabled
0/2	Disabled
0/3	Disabled
0/4	Disabled
0/5	Disabled
0/6	Disabled
0/7	Disabled
0/8	Disabled
0/9	Disabled
0/10	Disabled

Figure 4.222 Switching > IGMP Snooping > Multicast Router Configuration

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. When configuring the IGMP snooping multicast router settings, this field identifies the interface(s) that are being configured.
Multicast Router	Indicates whether the interface is enabled or disabled as a multicast router interface.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.

4.4.9.6 Multicast Router VLAN Status

If a multicast router is attached to the switch, its existence can be learned dynamically. You can also statically configure one or more VLANs on each interface to act as a multicast router interface, which is an interface that faces a multicast router or IGMP querier and receives multicast traffic.

To access this page, click **Switching > IGMP Snooping > Multicast Router VLAN Status**.




Figure 4.223 Switching > IGMP Snooping > Multicast Router VLAN Status

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. Only interfaces that are configured with multicast router VLANs appear in the table.
VLAN IDs	The ID of the VLAN configured as enabled for multicast routing on the associated interface.
Refresh	Click Refresh to update the screen.
Add	Click Add to enable IGMP snooping on a VLAN. The Multicast Router VLAN Configuration Menu displays. Click a VLAN ID to select it (or CTRL + click to select multiple VLAN IDs). Click the appropriate arrow to move the selected VLAN ID or VLAN IDs to the Configured VLAN IDs window. Click Submit to save the values and update the screen.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to disable IGMP snooping for the selected entries.

4.4.9.7 Multicast Router VLAN Configuration

Use the IGMP Snooping Multicast Router VLAN Configuration page to enable or disable specific VLANs as multicast router interfaces for a physical port or LAG. A multicast router interface faces a multicast router or IGMP querier and receives multicast traffic.

To access this page, click **Switching > IGMP Snooping > Multicast Router VLAN Configuration**.




Figure 4.224 Switching > IGMP Snooping > Multicast Router VLAN Configuration

The following table describes the items in the previous figure.

Item	Description
Interface	Click the drop-down menu to select the port or LAG on which to enable or disable a VLAN multicast routing interface.

Item	Description
VLAN IDs	The VLANs configured on the system that are not currently enabled as multicast router interfaces on the selected port or LAG. To enable a VLAN as a multicast router interface, click the VLAN ID to select it (or CTRL + click to select multiple VLAN IDs). Then, click the appropriate arrow to move the selected VLAN or VLANs to the Configured VLAN IDs window.
Configured VLAN IDs	The VLANs that are enabled as multicast router interfaces on the selected port or LAG. To disable a VLAN as a multicast router interface, click the VLAN ID to select it (or CTRL + click to select multiple VLAN IDs). Then, click the appropriate arrow to move the selected VLAN or VLANs to the VLAN IDs window.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.

4.4.10 IGMP Snooping Querier

4.4.10.1 Configuration

Use the IGMP Snooping Querier Configuration page to configure the global IGMP snooping querier settings on the device. IGMP snooping requires that one central switch or router periodically query all end-devices on the network to announce their multicast memberships. This central device is the IGMP querier. When layer 3 IP multicast routing protocols are enabled in a network with IP multicast routing, the IP multicast router acts as the IGMP querier. However, if the IP-multicast traffic in a VLAN needs to be layer 2 switched only, an IP-multicast router is not required. The IGMP snooping querier can perform the IGMP snooping functions on the VLAN.

To access this page, click **Switching > IGMP Snooping Querier > Configuration**.

Figure 4.225 Switching > IGMP Snooping Querier > Configuration

The following table describes the items in the previous figure.

Item	Description
Admin Mode	The administrative mode for the IGMP snooping querier on the device. When enabled, the IGMP snooping querier sends out periodic IGMP queries that trigger IGMP report messages from the switches that want to receive IP multicast traffic. The IGMP snooping feature listens to these IGMP reports to establish appropriate forwarding.
IP Address	The snooping querier address to be used as source address in periodic IGMP queries. This address is used when no IP address is configured on the VLAN on which the query is being sent.
IGMP Version	The IGMP protocol version used in periodic IGMP queries.
Query Interval (Seconds)	The amount of time the IGMP snooping querier on the device should wait between sending periodic IGMP queries.
Querier Expiry Interval (Seconds)	The amount of time the device remains in non-querier mode after it has discovered that there is a multicast querier in the network.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.4.10.2 VLAN Configuration

Use the IGMP Snooping Querier VLAN Configuration page to enable the IGMP snooping querier feature on one or more VLANs and to configure per-VLAN IGMP snooping querier settings. Only VLANs that have the IGMP snooping querier feature enabled appear in the table.

To access this page, click **Switching > IGMP Snooping Querier > VLAN Configuration**.




Figure 4.226 Switching > IGMP Snooping Querier > VLAN Configuration

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The VLAN on which the IGMP snooping querier is enabled. When enabling the IGMP snooping querier on a VLAN, use this menu to select the desired VLAN. Only VLANs that have been configured on the system and are not already enabled for the IGMP snooping querier appear in the menu. When modifying IGMP snooping querier settings, this field identifies the VLAN that is being configured.
Querier Election Participation	The participation mode for the IGMP snooping querier election process: <ul style="list-style-type: none"> ■ Enabled: The IGMP snooping querier on this VLAN participates in the querier election process when it discovers the presence of another querier in the VLAN. If the snooping querier finds that the other querier source IP address is lower than its own address, it stops sending periodic queries. If the snooping querier wins the election (because it has the lowest IP address), then it continues sending periodic queries. ■ Disabled: When the IGMP snooping querier on this VLAN sees other queriers of the same version in the VLAN, the snooping querier moves to the non-querier state and stops sending periodic queries.
Querier VLAN IP Address	The IGMP snooping querier address the VLAN uses as the source IP address in periodic IGMP queries sent on the VLAN. If this value is not configured, the VLAN uses the global IGMP snooping querier IP address.
Refresh	Click Refresh to update the screen.
Add	Click Add to enable the IGMP snooping querier feature on a VLAN. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to disable the IGMP snooping querier feature for the selected entries.

To enable the IGMP snooping querier feature on a VLAN:

Click **Switching > IGMP Snooping Querier > VLAN Configuration > Add**.




Figure 4.227 Switching > IGMP Snooping Querier > VLAN Configuration > Add

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The VLAN on which the IGMP snooping querier is enabled. When enabling the IGMP snooping querier on a VLAN, use this menu to select the desired VLAN. Only VLANs that have been configured on the system and are not already enabled for the IGMP snooping querier appear in the menu. When modifying IGMP snooping querier settings, this field identifies the VLAN that is being configured.
Querier Election Participation	The participation mode for the IGMP snooping querier election process: <ul style="list-style-type: none"> ■ Enabled: The IGMP snooping querier on this VLAN participates in the querier election process when it discovers the presence of another querier in the VLAN. If the snooping querier finds that the other querier source IP address is lower than its own address, it stops sending periodic queries. If the snooping querier wins the election (because it has the lowest IP address), then it continues sending periodic queries. ■ Disabled: When the IGMP snooping querier on this VLAN sees other queriers of the same version in the VLAN, the snooping querier moves to the non-querier state and stops sending periodic queries.
Querier VLAN IP Address	The IGMP snooping querier address the VLAN uses as the source IP address in periodic IGMP queries sent on the VLAN. If this value is not configured, the VLAN uses the global IGMP snooping querier IP address.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.4.10.3 VLAN Status

Use the IGMP Snooping Querier VLAN Status page to view information about the IGMP snooping querier status for all VLANs that have the snooping querier enabled.

To access this page, click **Switching > IGMP Snooping Querier > VLAN Status**.




Figure 4.228 Switching > IGMP Snooping Querier > VLAN Status

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The VLAN associated with the rest of the data in the row. The table includes only VLANs that have the snooping querier enabled.

Item	Description
State	The operational state of the IGMP snooping querier on the VLAN, which is one of the following: <ul style="list-style-type: none"> ■ Querier: The snooping switch is the querier in the VLAN. The snooping switch will send out periodic queries with a time interval equal to the configured querier query interval. If the snooping switch sees a better querier (numerically lower) in the VLAN, it moves to non-querier mode. ■ Non-Querier: The snooping switch is in non-querier mode in the VLAN. If the querier expiry interval timer expires, the snooping switch moves into querier mode. ■ Disabled: The snooping querier is not operational on the VLAN. The snooping querier moves to the disabled mode when IGMP snooping is not operational on the VLAN, when the querier address is not configured, or the network management address is not configured.
Version	The operational IGMP protocol version of the querier.
Last IP Address	The IP address of the last querier from which a query was snooped on the VLAN.
Last Version	The IGMP protocol version of the last querier from which a query was snooped on the VLAN.
Max Response Time (Seconds)	The maximum response time to be used in the queries that are sent by the snooping querier.
Refresh	Click Refresh to update the screen.

4.4.11 MLD Snooping

4.4.11.1 Configuration

Use the MLD Snooping Configuration and Status page to enable Multicast Listener Discovery (MLD) snooping on the device and to view global status information. In IPv4, Layer 2 switches can use IGMP snooping to limit the flooding of multicast traffic by dynamically configuring Layer 2 interfaces so that multicast traffic is forwarded to only those interfaces associated with IP multicast address. In IPv6 networks, MLD snooping performs a similar function. With MLD snooping, IPv6 multicast data is selectively forwarded to a list of ports intended to receive the data (instead of being flooded to all of the ports in a VLAN). This list is constructed by snooping IPv6 multicast control packets.

To access this page, click **Switching > MLD Snooping > Configuration**.

Figure 4.229 Switching > MLD Snooping > Configuration

The following table describes the items in the previous figure.

Item	Description
MLD Snooping Admin Mode	The administrative mode of MLD snooping on the device.
Multicast Control Frame Count	The number of multicast control frames that have been processed by the CPU.

Item	Description
Interfaces Enabled for MLD Snooping	One or more interfaces on which MLD snooping is administratively enabled. MLD snooping must be enabled globally and on an interface for the interface to be able to snoop MLD packets to determine which segments should receive multicast packets directed to the group address.
VLANs Enabled for MLD Snooping	One or more VLANs on which MLD snooping is administratively enabled.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.4.11.2 Interface Configuration

Use the MLD Snooping Interface Configuration page to configure MLD snooping settings on specific interfaces.

To access this page, click **Switching > MLD Snooping > Interface Configuration**.

The screenshot shows a web interface for configuring MLD snooping. At the top, it says 'Display 10 rows' and 'Showing 1 to 10 of 42 entries'. Below this is a table with columns: Interface, Admin Mode, Group Membership Interval, Max Response Time, Multicast Router Expiration Time, and Fast Leave Admin Mode. The table lists 10 interfaces (0/1 to 0/10), all with Admin Mode set to 'Disabled', Group Membership Interval of 260, Max Response Time of 10, Multicast Router Expiration Time of 0, and Fast Leave Admin Mode of 'Disabled'. At the bottom of the table, there are navigation buttons: 'First', 'Previous', '1', '2', '3', '4', '5', 'Next', and 'Last'. Below the table are two buttons: 'Refresh' and 'Edit'.

Figure 4.230 Switching > MLD Snooping > Interface Configuration

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. When configuring MLD snooping settings, this field identifies each interface that is being configured.
Admin Mode	The administrative mode of MLD snooping on the interface. MLD snooping must be enabled globally and on an interface for the interface to be able to snoop MLD packets to determine which segments should receive multicast packets directed to the group address.
Group Membership Interval	The number of seconds the interface should wait for a report for a particular group on the interface before the MLD snooping feature deletes the interface from the group.
Max Response Time	The number of seconds the interface should wait after sending a query if it does not receive a report for a particular group. The specified value should be less than the Group Membership Interval.
Multicast Router Expiration Time	The number of seconds the interface should wait to receive a query before it is removed from the list of interfaces with multicast routers attached.
Fast Leave Admin Mode	The administrative mode of Fast Leave on the interface. If Fast Leave is enabled, the interface can be immediately removed from the Layer 2 forwarding table entry upon receiving an MLD leave message for a multicast group without first sending out MAC-based general queries.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.

4.4.11.3 Source Specific Multicast

The MLD Snooping Source Specific Multicast page displays Source Specific Multicast (SSM) information learned by snooping MLDv2 reports. MLDv2 includes support for SSM, in which a receiver can request to receive multicast packets from one or more specific source address or from all addresses except one or more specified source addresses. If a host sends an MLDv2 report, the MLD snooping feature records the information and adds an entry to the table on this page.

To access this page, click **Switching > MLD Snooping > Source Specific Multicast**.




Figure 4.231 Switching > MLD Snooping > Source Specific Multicast

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The VLAN on which the MLDv2 report is received.
Group	The IPv6 multicast group address of the multicast group the host belongs to.
Interface	The interface on which the MLD v2 report is received.
Reporter	The IPv6 address of the host that sent the MLDv2 report.
Source Filter Mode	The source filter mode for the specified group, which is one of the following: <ul style="list-style-type: none"> ■ Include: The receiver has expressed interest in receiving multicast traffic for the multicast group from the source or sources in the Source Address List. ■ Exclude: The receiver has expressed interest in receiving multicast traffic for the multicast group from any source except the source or sources in the Source Address List.
Source Address List	The source IPv6 address or addresses for which source filtering is requested.
Refresh	Click Refresh to update the screen.

4.4.11.4 VLAN Status

Use the MLD Snooping VLAN Status page to enable or disable MLD snooping on system VLANs and to view and configure per-VLAN MLD snooping settings. Only VLANs that are enabled for MLD snooping appear in the table.

To access this page, click **Switching > MLD Snooping > VLAN Status**.




Figure 4.232 Switching > MLD Snooping > VLAN Status

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The VLAN associated with the rest of the data in the row. When enabling MLD snooping on a VLAN, use this menu to select the desired VLAN. Only VLANs that have been configured on the system and are not already enabled for MLD snooping appear in the menu. When modifying MLD snooping settings, this field identifies the VLAN that is being configured.
Admin Mode	The administrative mode of MLD snooping on the VLAN. MLD snooping must be enabled globally and on a VLAN for the VLAN to be able to snoop MLD packets and determine which network segments should receive multicast packets directed to the group address.
Group Membership Interval	The number of seconds the VLAN should wait for a report for a particular group on the VLAN before the MLD snooping feature deletes the VLAN from the group.
Max Response Time	The number of seconds the VLAN should wait after sending a query if does not receive a report for a particular group. The specified value should be less than the Group Membership Interval.
Multicast Router Expiration Time	The number of seconds the VLAN should wait to receive a query before it is removed from the list of VLANs with multicast routers attached.
Fast Leave Admin Mode	The administrative mode of Fast Leave on the VLAN. If Fast Leave is enabled, the VLAN can be immediately removed from the Layer 2 forwarding table entry upon receiving an MLD leave message for a multicast group without first sending out MAC-based general queries.
Refresh	Click Refresh to update the screen.
Add	Click Add to enable MLD snooping on a VLAN. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to disable MLD snooping for the selected entries.

To enable MLD snooping on a VLAN:

Click **Switching > MLD Snooping > VLAN Status > Add**.



Figure 4.233 Switching > MLD Snooping > VLAN Status > Add

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The VLAN associated with the rest of the data in the row. When enabling MLD snooping on a VLAN, use this menu to select the desired VLAN. Only VLANs that have been configured on the system and are not already enabled for MLD snooping appear in the menu. When modifying MLD snooping settings, this field identifies the VLAN that is being configured.

Item	Description
Group Membership Interval (Seconds)	The number of seconds the VLAN should wait for a report for a particular group on the VLAN before the MLD snooping feature deletes the VLAN from the group.
Max Response Time (Seconds)	The number of seconds the VLAN should wait after sending a query if does not receive a report for a particular group. The specified value should be less than the Group Membership Interval.
Multicast Router Expiration Time (Seconds)	The number of seconds the VLAN should wait to receive a query before it is removed from the list of VLANs with multicast routers attached.
Fast Leave Admin Mode	The administrative mode of Fast Leave on the VLAN. If Fast Leave is enabled, the VLAN can be immediately removed from the Layer 2 forwarding table entry upon receiving an MLD leave message for a multicast group without first sending out MAC-based general queries.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.4.11.5 Multicast Router Configuration

Use the MLD Snooping Multicast Router Configuration page to manually configure an interface as a static MLD snooping multicast router interface. If a multicast router is attached to the switch, its existence can be learned dynamically. You can also statically configure an interface as a multicast router interface, which is an interface that faces a multicast router or MLD querier and receives multicast traffic.

To access this page, click **Switching > MLD Snooping > Multicast Router Configuration**.

The screenshot shows a web interface for configuring Multicast Router settings. At the top, it says 'Display 10 rows' and 'Showing 1 to 10 of 42 entries'. There is a search filter box. The table has two columns: 'Interface' and 'Multicast Router'. The 'Interface' column lists interfaces from 0/1 to 0/10. The 'Multicast Router' column shows 'Disabled' for all listed interfaces. At the bottom of the table, there are navigation buttons: 'First', 'Previous', '1', '2', '3', '4', '5', 'Next', 'Last'. Below the table are 'Refresh' and 'Edit' buttons.

Interface	Multicast Router
0/1	Disabled
0/2	Disabled
0/3	Disabled
0/4	Disabled
0/5	Disabled
0/6	Disabled
0/7	Disabled
0/8	Disabled
0/9	Disabled
0/10	Disabled

Figure 4.234 Switching > MLD Snooping > Multicast Router Configuration

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. When configuring the MLD snooping multicast router settings, this field identifies each interface that is being configured.
Multicast Router	Indicates whether the interface is enabled or disabled as a multicast router interface.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.

4.4.11.6 Multicast Router VLAN Status

Use the MLD Snooping Multicast Router VLAN Status page to enable or disable specific VLANs as static multicast router interfaces for a physical port or LAG and to view the multicast router VLAN status for each interface. A multicast router interface faces a multicast router or MLD querier and receives multicast traffic. If a multicast router is attached to the switch, its existence can be learned dynamically. You can also statically configure one or more VLANs on each interface to act as multicast router interfaces.

To access this page, click **Switching > MLD Snooping > Multicast Router VLAN Status**.




Figure 4.235 Switching > MLD Snooping > Multicast Router VLAN Status

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. Only interfaces that are configured with multicast router VLANs appear in the table. When adding multicast router VLAN information for an interface, use the Interface menu to select the interface on which to enable one or more multicast router VLAN interfaces. When editing multicast router VLAN information, this field shows the interface that is being configured.
VLAN ID	The ID of each VLAN configured as enabled as a multicast router interface on the associated interface. When changing the multicast routing VLAN interfaces that are associated with an interface, click the VLAN ID to select it (or CTRL + click to select multiple VLAN IDs).
Refresh	Click Refresh to update the screen.
Add	Click Add to enable VLANs as multicast router interfaces on a port or LAG. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to disable all VLAN multicast routing interfaces for the selected entries.

To enable VLANs as multicast router interfaces on a port or LAG:

Click **Switching > MLD Snooping > Multicast Router VLAN Status > Add**.

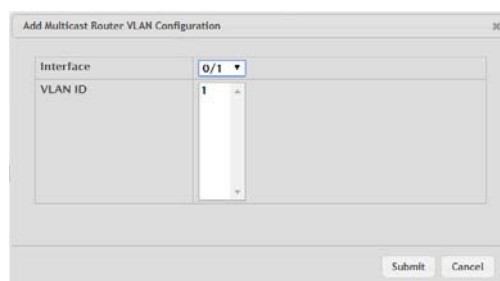


Figure 4.236 Switching > MLD Snooping > Multicast Router VLAN Status > Add

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. Only interfaces that are configured with multicast router VLANs appear in the table. When adding multicast router VLAN information for an interface, use the Interface menu to select the interface on which to enable one or more multicast router VLAN interfaces. When editing multicast router VLAN information, this field shows the interface that is being configured.
VLAN ID	The ID of each VLAN configured as enabled as a multicast router interface on the associated interface. When changing the multicast routing VLAN interfaces that are associated with an interface, click the VLAN ID to select it (or CTRL + click to select multiple VLAN IDs).
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.4.12 MLD Snooping Querier

4.4.12.1 Configuration

Use the MLD Snooping Querier Configuration page to configure the global MLD snooping querier settings on the device. MLD snooping requires that one central switch or router periodically query all end-devices on the network to announce their multicast memberships. This central device is the MLD querier. When Layer 3 IP multicast routing protocols are enabled in a network with IP multicast routing, the IP multicast router acts as the MLD querier. However, if the IP-multicast traffic in a VLAN needs to be Layer 2 switched only, an IP-multicast router is not required. The MLD snooping querier can perform the MLD snooping functions on the VLAN.

To access this page, click **Switching > MLD Snooping Querier > Configuration**.

Figure 4.237 Switching > MLD Snooping Querier > Configuration

The following table describes the items in the previous figure.

Item	Description
Admin Mode	The administrative mode for the MLD snooping querier on the device. When enabled, the MLD snooping querier sends out periodic MLD queries that trigger MLD report messages from the switches that want to receive IP multicast traffic. The MLD snooping feature listens to these MLD reports to establish appropriate forwarding.
IPv6 Address	The snooping querier unicast link-local IPv6 address to be used as the source address in periodic MLD queries. This address is used when no IPv6 address is configured on the VLAN on which the query is being sent.
MLD Version	The MLD protocol version used in periodic MLD queries.
Query Interval (Seconds)	The amount of time the MLD snooping querier should wait between sending periodic MLD queries.
Querier Expiry Interval (Seconds)	The amount of time the device remains in non-querier mode after it has discovered that there is a multicast querier in the network.
Submit	Click Submit to save the values and update the screen.

Item	Description
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.4.12.2 VLAN Configuration

Use the MLD Snooping Querier VLAN Configuration page to enable the MLD snooping querier feature on one or more VLANs and to configure per-VLAN MLD snooping querier settings. Only VLANs that have the MLD snooping querier feature enabled appear in the table.

To access this page, click **Switching > MLD Snooping Querier > VLAN Configuration**.




Figure 4.238 Switching > MLD Snooping Querier > VLAN Configuration

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The VLAN on which the MLD snooping querier is enabled. When enabling the MLD snooping querier on a VLAN, use this menu to select the desired VLAN. Only VLANs that have been configured on the system and are not already enabled for the MLD snooping querier appear in the menu. When modifying MLD snooping querier settings, this field identifies the VLAN that is being configured.
Querier Election Participation	The participation mode for the MLD snooping querier election process: <ul style="list-style-type: none"> ■ Enabled: The MLD snooping querier on this VLAN participates in the querier election process when it discovers the presence of another querier in the VLAN. If the snooping querier finds that the other querier source IPv6 address is lower than its own address, it stops sending periodic queries. If the snooping querier wins the election (because it has the lowest IPv6 address), then it continues sending periodic queries. ■ Disabled: When the MLD snooping querier on this VLAN sees other queriers of the same version in the VLAN, the snooping querier moves to the non-querier state and stops sending periodic queries.
Querier VLAN IPv6 Address	The MLD snooping querier unicast link-local IPv6 address the VLAN uses as the source address in periodic MLD queries sent on the VLAN. If this value is not configured, the VLAN uses the global MLD snooping querier IPv6 address.
Refresh	Click Refresh to update the screen.
Add	Click Add to enable the MLD snooping querier feature on a VLAN. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to disable the MLD snooping querier feature for the selected entries.

To enable the MLD snooping querier feature on a VLAN:

Click **Switching > MLD Snooping Querier > VLAN Configuration > Add**.




Figure 4.239 Switching > MLD Snooping Querier > VLAN Configuration > Add

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The VLAN on which the MLD snooping querier is enabled. When enabling the MLD snooping querier on a VLAN, use this menu to select the desired VLAN. Only VLANs that have been configured on the system and are not already enabled for the MLD snooping querier appear in the menu. When modifying MLD snooping querier settings, this field identifies the VLAN that is being configured.
Querier Election Participation	The participation mode for the MLD snooping querier election process: <ul style="list-style-type: none"> ■ Enabled: The MLD snooping querier on this VLAN participates in the querier election process when it discovers the presence of another querier in the VLAN. If the snooping querier finds that the other querier source IPv6 address is lower than its own address, it stops sending periodic queries. If the snooping querier wins the election (because it has the lowest IPv6 address), then it continues sending periodic queries. ■ Disabled: When the MLD snooping querier on this VLAN sees other queriers of the same version in the VLAN, the snooping querier moves to the non-querier state and stops sending periodic queries.
Querier VLAN IPv6 Address	The MLD snooping querier unicast link-local IPv6 address the VLAN uses as the source address in periodic MLD queries sent on the VLAN. If this value is not configured, the VLAN uses the global MLD snooping querier IPv6 address.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.4.12.3 VLAN Status

Use the MLD Snooping Querier VLAN Status page to view information about the MLD snooping querier status for all VLANs that have the snooping querier enabled.

To access this page, click **Switching > MLD Snooping Querier > VLAN Status**.




Figure 4.240 Switching > MLD Snooping Querier > VLAN Status

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The VLAN associated with the rest of the data in the row. The table includes only VLANs that have the snooping querier enabled.

Item	Description
State	The operational state of the MLD Snooping Querier on a VLAN, which is one of the following: <ul style="list-style-type: none"> ■ Querier: The snooping switch is the querier in the VLAN. The snooping switch will send out periodic queries with a time interval equal to the configured querier query interval. If the snooping switch sees a better querier (numerically lower) in the VLAN, it moves to non-querier mode. ■ Non-Querier: The snooping switch is in non-querier mode in the VLAN. If the querier expiry interval timer expires, the snooping switch moves into querier mode. ■ Disabled: The snooping querier is not operational on the VLAN. The snooping querier moves to the disabled mode when MLD snooping is not operational on the VLAN, when the querier address is not configured, or the network management address is not configured.
Version	The operational MLD protocol version of the querier.
Last IPv6 Address	The IPv6 address of the last querier from which a query was snooped on the VLAN.
Last Version	The MLD protocol version of the last querier from which a query was snooped on the VLAN.
Max Response Time (Seconds)	The maximum response time to be used in the queries that are sent by the snooping querier.
Refresh	Click Refresh to update the screen.

4.4.13 Multicast Forwarding Database

The Layer 2 Multicast Forwarding Database (MFDB) is used by the switch to make forwarding decisions for packets that arrive with a multicast destination MAC address. By limiting multicasts to only certain ports in the switch, traffic is prevented from going to parts of the network where that traffic is unnecessary.

When a packet enters the switch, the destination MAC address is combined with the VLAN ID and a search is performed in the Layer 2 Multicast Forwarding Database. If no match is found, the packet is either flooded to all ports in the VLAN or discarded, depending on the switch configuration. If a match is found, the packet is forwarded only to the ports that are members of that multicast group.

4.4.13.1 Summary

The Multicast Forwarding Database Summary page displays the entries in the multicast forwarding database (MFDB) on the device. The MFDB holds the port membership information for all active multicast address entries and is used to make forwarding decisions for frames that arrive with a multicast destination MAC address. The key for an entry consists of a VLAN ID and MAC address pair. Entries may contain data for more than one protocol.

To access this page, click **Switching > Multicast Forwarding Database > Summary**.




Figure 4.241 Switching > Multicast Forwarding Database > Summary

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The VLAN ID associated with the entry in the MFDB.
MAC Address	The multicast MAC address that has been added to the MFDB.
Component	The feature on the device that was responsible for adding the entry to the multicast forwarding database, which is one of the following: <ul style="list-style-type: none"> ■ IGMP Snooping: A layer 2 feature that allows the device to dynamically add or remove ports from IPv4 multicast groups by listening to IGMP join and leave requests. ■ MLD Snooping: A layer 2 feature that allows the device to dynamically add or remove ports from IPv6 multicast groups by listening to MLD join and leave requests. ■ GMRP: Generic Address Resolution Protocol (GARP) Multicast Registration Protocol, which helps control the flooding of multicast traffic by keeping track of group membership information. ■ Static Filtering: A static MAC filter that was manually added to the address table by an administrator.
Type	The type of entry, which is one of the following: <ul style="list-style-type: none"> ■ Static: The entry has been manually added to the MFDB by an administrator. ■ Dynamic: The entry has been added to the MFDB as a result of a learning process or protocol.
Description	A text description of this multicast table entry.
Interface(s)	The list of interfaces that will forward or filter traffic sent to the multicast MAC address.
Forwarding Interface(s)	The list of forwarding interfaces. This list does not include any interfaces that are listed as static filtering interfaces.
Refresh	Click Refresh to update the screen.

4.4.13.2 GMRP

Use the Multicast Forwarding Database GMRP Table page to display the entries in the multicast forwarding database (MFDB) that were added by using the GARP Multicast Registration Protocol (GMRP).

To access this page, click **Switching > Multicast Forwarding Database > GMRP**.




Figure 4.242 Switching > Multicast Forwarding Database > GMRP

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The VLAN ID associated with the entry in the MFDB.
MAC Address	The multicast MAC address associated with the entry in the MFDB.
Type	The type of entry, which is one of the following: <ul style="list-style-type: none"> ■ Static: The entry has been manually added to the MFDB by an administrator. ■ Dynamic: The entry has been added to the MFDB as a result of a learning process or protocol. Entries that appear on this page have been added by using GARP.

Item	Description
Description	A text description of this multicast table entry.
Interface(s)	The list of interfaces that will forward or filter traffic sent to the multicast MAC address.
Refresh	Click Refresh to update the screen.

4.4.13.3 IGMP Snooping

The Multicast Forwarding Database IGMP Snooping Table page displays the entries in the multicast forwarding database (MFDB) that were added because they were discovered by the IGMP snooping feature. IGMP snooping allows the device to dynamically add or remove ports from IPv4 multicast groups by listening to IGMP join and leave requests.

To access this page, click **Switching > Multicast Forwarding Database > IGMP Snooping**.




Figure 4.243 Switching > Multicast Forwarding Database > IGMP Snooping

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The VLAN ID associated with the entry in the MFDB.
MAC Address	The multicast MAC address associated with the entry in the MFDB.
Type	The type of entry, which is one of the following: <ul style="list-style-type: none"> ■ Static: The entry has been manually added to the MFDB by an administrator. ■ Dynamic: The entry has been added to the MFDB as a result of a learning process or protocol. Entries that appear on this page have been learned by examining IGMP messages.
Description	A text description of this multicast table entry.
Interface(s)	The list of interfaces that will forward or filter traffic sent to the multicast MAC address.
Refresh	Click Refresh to update the screen.
Clear Entries	Click Clear Entries to remove all IGMP snooping entries from the MFDB table.

4.4.13.4 Source Specific Multicast

The MFDB Source Specific Multicast Table page displays the entries in the multicast forwarding database (MFDB) for source specific multicast, that were added because they were discovered by the IGMP Snooping or MLD Snooping feature.

To access this page, click **Switching > Multicast Forwarding Database > Source Specific Multicast**.



Figure 4.244 Switching > Multicast Forwarding Database > Source Specific Multicast

The following table describes the items in the previous figure.

Item	Description
Type	Type of snooping. The values can be either IGMP Snooping or MLD Snooping.
VLAN ID	VLAN on which the entry has been learned.
Group	The multicast group address.
Source	The source address.
Source Filter Mode	The source filter mode(Include/Exclude) for the specified group.
Interface(s)	Specifies the list of interfaces on which a incoming packet is forwarded.
Refresh	Click Refresh to update the screen.

4.4.13.5 Source Specific Multicast Status

The MFDB Source Specific Multicast Status page displays information about the source specific multicast forwarding database (SSMFDB).

To access this page, click **Switching > Multicast Forwarding Database > Source Specific Multicast Status**.

IGMP Snooping	
Total Entries	16
Peak Entries	0
Current Entries	0
MLD Snooping	
Total Entries	8
Peak Entries	0
Current Entries	0

Refresh

Figure 4.245 Switching > Multicast Forwarding Database > Source Specific Multicast Status

The following table describes the items in the previous figure.

Item	Description
IGMP Snooping	
Total Entries	The total number of entries that can possibly be in IGMP snooping's SSMFDB.
Peak Entries	The largest number of entries that have been present in the IGMP snooping's SSMFDB.
Current Entries	The current number of entries in the IGMP snooping's SSMFDB.
MLD Snooping	
Total Entries	The total number of entries that can possibly be in MLD snooping's SSMFDB.
Peak Entries	The largest number of entries that have been present in the MLD snooping's SSMFDB.
Current Entries	The current number of entries in the MLD snooping's SSMFDB.
Refresh	Click Refresh to update the screen.

4.4.13.6 Statistics

The Multicast Forwarding Database Statistics page displays statistical information about the multicast forwarding database (MFDB).

To access this page, click **Switching > Multicast Forwarding Database > Statistics**.

MFDB Max Table Entries	1024
MFDB Most Entries Since Last Reset	0
MFDB Current Entries	0

Figure 4.246 Switching > Multicast Forwarding Database > Statistics

The following table describes the items in the previous figure.

Item	Description
MFDB Max Table Entries	The maximum number of entries that the multicast forwarding database can hold.
MFDB Most Entries Since Last Reset	The largest number of entries that have been present in the multicast forwarding database since the device was last reset. This value is also known as the MFDB high-water mark.
MFDB Current Entries	The current number of entries in the multicast forwarding database.
Refresh	Click Refresh to update the screen.

4.4.14 MVR

Multicast VLAN Registration (MVR) allows the switch to listen to the Internet Group Management Protocol (IGMP) frames. Both protocols operate independently from each other and can be enabled on the switch interfaces. In such case, MVR listens to the Join and Report messages only for the statically configured groups. All other groups are managed by IGMP snooping. MVR uses the multicast VLAN, a dedicated VLAN used to transfer multicast traffic over the network avoiding duplication of multicast streams for clients in different VLANs.

4.4.14.1 Global

Use the MVR Global Configuration page to view and configure the global settings for MVR.

To access this page, click **Switching > MVR > Global**.

Admin Mode	<input type="radio"/> Enable <input checked="" type="radio"/> Disable
MVR Mode	<input checked="" type="radio"/> Compatible <input type="radio"/> Dynamic
Multicast VLAN	<input type="text" value="1"/> (1 to 4093)
Maximum Multicast Groups	256
Current Multicast Groups	0
Query Response Time (Tenths of Seconds)	<input type="text" value="5"/> (1 to 100)

Figure 4.247 Switching > MVR > Global

The following table describes the items in the previous figure.

Item	Description
Admin Mode	The administrative mode of MVR on the device.

Item	Description
MVR Mode	<p>The MVR learning mode, which can be one of the following:</p> <ul style="list-style-type: none"> Compatible: MVR does not learn source ports membership, instead all source ports are members of all groups by default. MVR does not forward IGMP Joins and Leaves from the hosts to the router. Dynamic: MVR learns source ports membership from IGMP queries. MVR forwards the IGMP Joins and Leaves from the hosts to the router. <p>The multicast traffic is forwarded only to the receiver ports that joined the group, either by IGMP Joins or MVR static configuration.</p>
Multicast VLAN	A dedicated VLAN used to transfer multicast traffic over the network avoiding duplication of multicast streams for clients in different VLANs.
Maximum Multicast Groups	The maximum number of membership groups that can be statically configured in the MVR database.
Current Multicast Groups	The current number of membership groups that are statically configured in the MVR database.
Query Response Time (Tenths of Seconds)	The maximum time to wait for an IGMP membership report on a receiver port before removing the port from the multicast group. The query time is specified in tenths of a second.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.4.14.2 Group

Use the MVR Group Status page to view or configure MVR groups. MVR maintains two types of group entries in its database, Static and Dynamic. Static entries are configured by the administrator and Dynamic entries are learned by MVR on the source ports.

To access this page, click **Switching > MVR > Group**.



Figure 4.248 Switching > MVR > Group

The following table describes the items in the previous figure.

Item	Description
Group	The multicast group address.
Status	<p>The status of the group, which can be one of the following:</p> <ul style="list-style-type: none"> Active: Group has one or more MVR ports participating. Inactive: Group has no MVR ports participating.
Members	The list of interfaces which participate in the MVR group. In the compatible mode, all source ports are members of all groups by default.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new group. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

To add a new group:

Click **Switching > MVR > Group > Add**.

Figure 4.249 Switching > MVR > Group > Add

The following table describes the items in the previous figure.

Item	Description
Group	The multicast group address.
Contiguous Group Count	Specify the desired number of groups to be created starting with the entered group address. The default contiguous group count is 1.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.4.14.3 Interface

Use the MVR Interface Status page to configure MVR settings on specific interfaces. To configure the settings for one or more interfaces, select each entry to modify and click Edit. The same MVR settings are applied to all selected interfaces.

To access this page, click **Switching > MVR > Interface**.

Figure 4.250 Switching > MVR > Interface

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. When configuring MVR settings, this field identifies the interface(s) that are being configured.
MVR Interface Mode	The administrative mode of MVR on the interface. MVR must be enabled globally and on an interface in order to listen to the Join and Report messages for the configured groups.
Type	The type of interface, which can be one of the following: <ul style="list-style-type: none"> ■ Source: The port where multicast traffic is flowing to. It must be a member of the multicast VLAN. ■ Receiver: The port where listening host is connected to the switch. It must not be a member of the multicast VLAN. ■ None: The port is not an MVR port.

Item	Description
Status	<p>The active state of the interface, which can be one of the following:</p> <ul style="list-style-type: none"> ■ Active: The port has link up and is in the forwarding state. ■ Inactive: The port may not have link up, not be in the forwarding state, or both. <p>The interface VLAN information is also displayed as part of the status and can be one of the following:</p> <ul style="list-style-type: none"> ■ In VLAN: Interface is a member of one or more VLANs. ■ Not In VLAN: Interface is not a member of any VLAN.
Immediate Leave	<p>The MVR immediate leave mode on the interface. It can only be configured on the receiver ports. MVR handles IGMP Leaves in Normal or Immediate leave mode. When a Leave message is received, in the normal mode a general IGMP query is sent from the switch to the receiver port, whereas in the immediate leave mode the switch is immediately reconfigured not to forward specific multicast stream to the receiver port. The immediate leave mode is used for ports where only one client may be connected.</p>
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.

4.4.14.4 Statistics

Use the MVR Statistics page to view statistical information about IGMP packets intercepted by MVR.

To access this page, click **Switching > MVR > Statistics**.

Statistics	Transmit	Receive
IGMP Queries	0	0
IGMPv1 Reports	0	0
IGMPv2 Reports	0	0
IGMP Leaves	0	0
Packet Failures	0	0

Figure 4.251 Switching > MVR > Statistics

The following table describes the items in the previous figure.

Item	Description
IGMP Queries	The total number of IGMP Queries successfully transmitted or received by the processor.
IGMPv1 Reports	The total number of IGMPv1 Reports successfully transmitted or received by the processor.
IGMPv2 Reports	The total number of IGMPv2 Reports successfully transmitted or received by the processor.
IGMP Leaves	The total number of IGMP Leaves successfully transmitted or received by the processor.
Packet Failures	The total number of packets which failed to get transmitted or received by the processor.
Refresh	Click Refresh to update the screen.

4.4.15 LLDP

The IEEE 802.1AB defined standard, Link Layer Discovery Protocol (LLDP), allows stations residing on an 802 LAN to advertise major capabilities and physical descriptions. This information is viewed by a network manager to identify system topology and detect bad configurations on the LAN.

LLDP is a one-way protocol; there are no request/response sequences. Information is advertised by stations implementing the transmit function, and is received and processed by stations implementing the receive function. The transmit and receive functions can be enabled/disabled separately per port. By default, both transmit and receive are disabled on all ports. The application is responsible for starting each transmit and receive state machine appropriately, based on the configured status and operational state of the port.

FASTPATH SMB allows LLDP to have multiple LLDP neighbors per interface. The number of such neighbors is limited by the memory constraints. A product-specific constant defines the maximum number of neighbors supported by the switch. There is no restriction on the number of neighbors supported on a per LLDP port. If all the remote entries on the switch are filled up, the new neighbors are ignored. In case of multiple VOIP devices on a single interface, the 802.1ab component sends the Voice VLAN configuration to all the VoIP devices.

4.4.15.1 Global

Use the LLDP Global Configuration page to specify LLDP parameters that are applied to the switch.

To access this page, click **Switching > LLDP > Global**.

Transmit Interval (Seconds)	30	(5 to 32768)
Transmit Hold Multiplier (Seconds)	4	(2 to 10)
Re-Initialization Delay (Seconds)	2	(1 to 10)
Notification Interval (Seconds)	5	(5 to 3600)

Submit Refresh Cancel

Figure 4.252 Switching > LLDP > Global

The following table describes the items in the previous figure.

Item	Description
Transmit Interval (Seconds)	The number of seconds between transmissions of LLDP advertisements.
Transmit Hold Multiplier (Seconds)	The Transmit Interval multiplier value, where $\text{Transmit Hold Multiplier} \times \text{Transmit Interval} = \text{the time to live (TTL) value the device advertises to neighbors}$.
Re-Initialization Delay (Seconds)	The number of seconds to wait before attempting to reinitialize LLDP on a port after the LLDP operating mode on the port changes.
Notification Interval (Seconds)	The minimum number of seconds to wait between transmissions of remote data change notifications to the SNMP trap receiver(s) configured on the device.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.4.15.2 Interface

Use the LLDP Interface Summary page to specify LLDP parameters that are applied to a specific interface.

To access this page, click **Switching > LLDP > Interface**.




Figure 4.253 Switching > LLDP > Interface

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. Only interfaces that have at least one LLDP setting enabled appear in the table. In the Add LLDP Interface window, use this field to select the interface with the LLDP settings to configure. In the Edit LLDP Interface window, this field identifies the interface that is being configured.
Port ID Subtype	The LLDP Port ID subtype of the interface, which is either MAC Address or Interface Name.
Link Status	The link status of the interface, which is either Up or Down. An interface that is down does not forward traffic.
Transmit	The LLDP advertise (transmit) mode on the interface. If the transmit mode is enabled, the interface sends LLDP Data Units (LLDPDUs) that advertise the mandatory TLVs and any optional TLVs that are enabled.
Receive	The LLDP receive mode on the interface. If the receive mode is enabled, the device can receive LLDPDUs from other devices.
Notify	The LLDP remote data change notification status on the interface. If the notify mode is enabled, the interface sends SNMP notifications when a link partner device is added or removed.
Optional TLV(s)	Indicates which optional LLDP TLV(s) are included in the LLDPDUs that the interface transmits: <ul style="list-style-type: none"> ■ 0: Port Description ■ 1: System Name ■ 2: System Description ■ 3: System Capabilities
Transmit Management Information	Indicates whether management address information for the local device is transmitted in LLDPDUs. Other remote managers can obtain information about the device by using its advertised management address.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new LLDP interface. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

To add a new LLDP interface:

Click **Switching > LLDP > Interface > Add**.

Figure 4.254 Switching > LLDP > Interface > Add

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. Only interfaces that have at least one LLDP setting enabled appear in the table. In the Add LLDP Interface window, use this field to select the interface with the LLDP settings to configure. In the Edit LLDP Interface window, this field identifies the interface that is being configured.
Port ID Subtype	The LLDP Port ID subtype of the interface, which is either MAC Address or Interface Name.
Transmit	The LLDP advertise (transmit) mode on the interface. If the transmit mode is enabled, the interface sends LLDP Data Units (LLDPDUs) that advertise the mandatory TLVs and any optional TLVs that are enabled.
Receive	The LLDP receive mode on the interface. If the receive mode is enabled, the device can receive LLDPDUs from other devices.
Notify	The LLDP remote data change notification status on the interface. If the notify mode is enabled, the interface sends SNMP notifications when a link partner device is added or removed.
Transmit Management Information	Indicates whether management address information for the local device is transmitted in LLDPDUs. Other remote managers can obtain information about the device by using its advertised management address.
Optional TLV(s)	
Port Description	Select this option to include the user-configured port description in the LLDPDU the interface transmits.
System Name	Select this option to include the user-configured system name in the LLDPDU the interface transmits. The system name is configured on the System Description page and is the SNMP server name for the device.
System Description	Select this option to include a description of the device in the LLDPDU the interface transmits. The description includes information about the product model and platform.
System Capabilities	Select this option to advertise the primary function(s) of the device in the LLDPDU the interface transmits.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.4.15.3 Local Devices

The LLDP Local Device Summary page displays summary information about the Link Layer Discovery Protocol (LLDP) data each interface advertises in the LLDP data units (LLDPDUs) it transmits. An interface appears in the table only if its LLDP transmit setting is enabled. To view additional LLDP information that the interface advertises, select the interface with the information to view and click Details.

To access this page, click **Switching > LLDP > Local Devices**.




Figure 4.255 Switching > LLDP > Local Devices

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the LLDP - 802.1AB data in the row. When viewing the details for an interface, this field identifies the interface that is being viewed.
Port ID	The port identifier, which is the physical address associated with the interface.
Port Description	A description of the port. An administrator can configure this information on the Port Description page.
Refresh	Click Refresh to update the screen.
Details	Click Details to open a window and display additional information.

4.4.15.4 Remote Devices

The LLDP Remote Device Summary page displays information about the remote devices the local system has learned about through the Link Layer Discovery Protocol (LLDP) data units received on its interfaces. The table lists all interfaces that are enabled to receive LLDP data from remote devices. However, information is available about remote devices only if the interface receives an LLDP data unit (LLDPDU) from a device. To view additional information about a remote device, select the interface that received the LLDP data and click Details.

To access this page, click **Switching > LLDP > Remote Devices**.




Figure 4.256 Switching > LLDP > Remote Devices

The following table describes the items in the previous figure.

Item	Description
Interface	The local interface that is enabled to receive LLDPDUs from remote devices.
Remote ID	The client identifier assigned to the remote system that sent the LLDPDU.
Chassis ID	The information the remote device sent as the Chassis ID TVL. This identifies the hardware platform for the remote system.
Port ID	The port on the remote system that transmitted the LLDP data.

Item	Description
System Name	The system name configured on the remote device.
Refresh	Click Refresh to update the screen.
Details	Click Details to open a window and display additional information.

4.4.15.5 Statistics

The LLDP Statistics page displays statistical information about the Link Layer Discovery Protocol (LLDP) Data Units (LLDPDUs) the interfaces on the local device have sent and received. The table that shows per-interface statistics contains entries only for interfaces that have at least one LLDP setting enabled.

To access this page, click **Switching > LLDP > Statistics**.

The screenshot shows the LLDP Statistics page. At the top, there is a summary table with the following data:

Last Update	0d:00:00:00
Total Inserts	0
Total Deletes	0
Total Drops	0
Total Ageouts	0

Below the summary table, there is a main table with the following headers: Interface, Transmit Total, Receive Total, Discards, Errors, Ageouts, TLV Discards, TLV Unknowns, TLV MED, TLV 802.1, TLV 802.3. The main table is currently empty, displaying "Table is Empty". There are navigation buttons (First, Previous, Next, Last) and "Refresh" and "Clear" buttons at the bottom of the table area.

Figure 4.257 Switching > LLDP > Statistics

The following table describes the items in the previous figure.

Item	Description
Last Update	The amount of time that has passed since an entry was created, modified, or deleted in the local database that maintains LLDP information received from remote systems.
Total Inserts	The number of times the complete set of information advertised by a particular MAC Service Access Point (MSAP) has been inserted into tables associated with the remote systems.
Total Deletes	The number of times the complete set of information advertised by a particular MSAP has been deleted from tables associated with the remote systems.
Total Drops	The number of times the complete set of information advertised by a particular MSAP could not be entered into tables associated with the remote systems because of insufficient resources.
Total Ageouts	The number of times the complete set of information advertised by a particular MSAP has been deleted from tables associated with the remote systems because the information timeliness interval has expired.
Interface	The interface associated with the rest of the data in the row.
Transmit Total	The number of LLDPDUs transmitted by the LLDP agent on the interface.
Receive Total	The number of valid LLDPDUs received by this interface while the LLDP agent is enabled.
Discards	The number of LLDP TLVs discarded for any reason by the LLDP agent on the interface.
Errors	The number of invalid LLDPDUs received by the LLDP agent on the interface while the LLDP agent is enabled.
Ageouts	The number of age-outs that have occurred on the interface. An age-out occurs the complete set of information advertised by a particular MSAP has been deleted from tables associated with the remote entries because the information timeliness interval had expired.

Item	Description
TLV Discards	The number of LLDP TLVs discarded for any reason by the LLDP agent on the interface.
TLV Unknowns	The number of LLDP TLVs received on the interface that were not recognized by the LLDP agent.
TLV MED	The total number of LLDP-MED TLVs received on the interface.
TLV 802.1	The total number of LLDP TLVs received on the interface which are of type 802.1.
TLV 802.3	The total number of LLDP TLVs received on the interface which are of type 802.3.
Refresh	Click Refresh to update the screen.
Clear	Click Clear to reset all LLDP statistics counters to zero.

4.4.16 LLDP-MED

The Link Layer Discovery Protocol-Media Endpoint Discovery (LLDP-MED) is an enhancement to LLDP that features:

- Auto-discovery of LAN policies (such as VLAN, Layer 2 Priority and DiffServ settings), enabling plug and play networking.
- Device location discovery for creation of location databases.
- Extended and automated power management of Power over Ethernet endpoints.
- Inventory management, enabling network administrators to track their network devices and determine their characteristics (manufacturer, software and hardware versions, serial/asset number).

4.4.16.1 Global

Use the LLDP-MED Global Configuration page to set global parameters for LLDP-MED operation.

To access this page, click **Switching > LLDP-MED > Global**.

Figure 4.258 Switching > LLDP-MED > Global

The following table describes the items in the previous figure.

Item	Description
Fast Start Repeat Count	The number of LLDP-MED Protocol Data Units (PDUs) that will be transmitted when the protocol is enabled.
Device Class	The device's MED Classification. The following three classifications represent the actual endpoints: <ul style="list-style-type: none"> ■ Class I Generic (for example, IP Communication Controller) ■ Class II Media (for example, Conference Bridge) ■ Class III Communication (for example, IP Telephone) The fourth device is Network Connectivity Device, which is typically a device such as a LAN switch or router, IEEE 802.1 bridge, or IEEE 802.11 wireless access point.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.4.16.2 Interface

Use the LLDP-MED Interface Summary page to enable LLDP-MED mode on an interface and to configure its properties.

To access this page, click **Switching > LLDP-MED > Interface**.

Interface	Link Status	MED Status	Notification Status	Operational Status	Transmit TLVs
0/1	Down	Disable	Disable	Disable	0, 1
0/2	Down	Disable	Disable	Disable	0, 1
0/3	Down	Disable	Disable	Disable	0, 1
0/4	Down	Disable	Disable	Disable	0, 1
0/5	Down	Disable	Disable	Disable	0, 1
0/6	Down	Disable	Disable	Disable	0, 1
0/7	Down	Disable	Disable	Disable	0, 1
0/8	Down	Disable	Disable	Disable	0, 1
0/9	Down	Disable	Disable	Disable	0, 1
0/10	Up	Disable	Disable	Disable	0, 1

Figure 4.259 Switching > LLDP-MED > Interface

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. When configuring LLDP-MED settings, this field identifies the interfaces that are being configured.
Link Status	The link status of the interface, which is either Up or Down. An interface that is down does not forward traffic.
MED Status / LLDP-MED Mode	The administrative status of LLDP-MED on the interface. When LLDP-MED is enabled, the transmit and receive function of LLDP is effectively enabled on the interface.
Notification Status / Configuration Notification Mode	Indicates whether LLDP-MED topology change notifications are enabled or disabled on the interface.
Operational Status	Indicates whether the interface will transmit TLVs.
Transmit TLVs	The LLDP-MED TLV(s) that the interface transmits: <ul style="list-style-type: none"> ■ Capabilities: 0 ■ Network Policy: 1
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new LLDP interface. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

To add a new LLDP interface:

Click **Switching > LLDP-MED > Interface > Add**.

The dialog box 'Add LLDP-MED Interface' contains the following fields and options:

- Interface:** A dropdown menu showing 'ge0/12'.
- LLDP-MED Mode:** Radio buttons for 'Enable' (selected) and 'Disable'.
- Configuration Notification Mode:** Radio buttons for 'Enable' (selected) and 'Disable'.
- Transmit TLVs:** A section with three checkboxes:
 - Capabilities:
 - Network Policy:
 - Extended PSE:

Buttons for 'Submit' and 'Cancel' are located at the bottom right.

Figure 4.260 Switching > LLDP-MED > Interface > Add

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. When configuring LLDP-MED settings, this field identifies the interfaces that are being configured.
MED Status / LLDP-MED Mode	The administrative status of LLDP-MED on the interface. When LLDP-MED is enabled, the transmit and receive function of LLDP is effectively enabled on the interface.
Notification Status / Configuration Notification Mode	Indicates whether LLDP-MED topology change notifications are enabled or disabled on the interface.
Transmit TLVs	The LLDP-MED TLV(s) that the interface transmits: <ul style="list-style-type: none"> ■ Capabilities: 0 ■ Network Policy: 1
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.4.16.3 Local Devices

The LLDP-MED Local Device Summary page displays information on LLDP-MED information advertised on the selected local interface.

To access this page, click **Switching > LLDP-MED > Local Devices**.




Figure 4.261 Switching > LLDP-MED > Local Devices

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. When viewing LLDP-MED details for an interface, this field identifies the interface that is being viewed.
Port ID	The MAC address of the interface. This is the MAC address that is advertised in LLDP-MED PDUs.
Refresh	Click Refresh to update the screen.
Details	Click Details to open a window and display additional information.

4.4.16.4 Remote Devices

The LLDP-MED Remote Device Summary page displays information about the remote devices the local system has learned about through the LLDP-MED data units received on its interfaces. Information is available about remote devices only if an interface receives an LLDP-MED data unit from a device. To view additional information about a remote device, select the interface that received the LLDP-MED data and click Details. The information below is organized according to the order in which the fields appear in the LLDP-MED Remote Device Information window.

To access this page, click **Switching > LLDP-MED > Remote Devices**.



Figure 4.262 Switching > LLDP-MED > Remote Devices

The following table describes the items in the previous figure.

Item	Description
Interface	The local interface that has received LLDP-MED data units from remote devices.
Remote ID	The client identifier assigned to the remote system that sent the LLDP-MED data unit.
Device Class	<p>The MED Classification advertised by the TLV from the remote device. The following three classifications represent the actual endpoints:</p> <ul style="list-style-type: none"> ■ Class I Generic (for example, IP Communication Controller) ■ Class II Media (for example, Conference Bridge) ■ Class III Communication (for example, IP Telephone) <p>The fourth device is Network Connectivity Device, which is typically a device such as a LAN switch or router, IEEE 802.1 bridge, or IEEE 802.11 wireless access point.</p>
Refresh	Click Refresh to update the screen.
Details	Click Details to open a window and display additional information.

4.4.17 Port Channel

4.4.17.1 Summary

Use the Port Channel Summary page to view and manage port channels on the device. Port channels, also known as Link Aggregation Groups (LAGs), allow one or more full-duplex Ethernet links of the same speed to be aggregated together. This allows the device to treat the port channel as a single, logical link. The primary purpose of a port channel is to increase the bandwidth between two devices. Port channels can also provide redundancy.

To access this page, click **Switching > Port Channel > Summary**.

Name	Type	Admin Mode	STP Mode	Link State	Link Trap	Members	Active Ports	Load Balance
ch1	Static	Enable	Enable	Down	Disable			Source/Destination MAC, VLAN, EtherType, Incoming Port
ch2	Static	Enable	Enable	Down	Disable			Source/Destination MAC, VLAN, EtherType, Incoming Port
ch3	Static	Enable	Enable	Down	Disable			Source/Destination MAC, VLAN, EtherType, Incoming Port
ch4	Static	Enable	Enable	Down	Disable			Source/Destination MAC, VLAN, EtherType, Incoming Port
ch5	Static	Enable	Enable	Down	Disable			Source/Destination MAC, VLAN, EtherType, Incoming Port
ch6	Static	Enable	Enable	Down	Disable			Source/Destination MAC, VLAN, EtherType, Incoming Port
ch7	Static	Enable	Enable	Down	Disable			Source/Destination MAC, VLAN, EtherType, Incoming Port
ch8	Static	Enable	Enable	Down	Disable			Source/Destination MAC, VLAN, EtherType, Incoming Port
ch9	Static	Enable	Enable	Down	Disable			Source/Destination MAC, VLAN, EtherType, Incoming Port
ch10	Static	Enable	Enable	Down	Disable			Source/Destination MAC, VLAN, EtherType, Incoming Port

Figure 4.263 Switching > Port Channel > Summary

The following table describes the items in the previous figure.

Item	Description
Name	A unique name to identify the port channel. Depending on the type of port channel, this name is automatically assigned by the system or can be configured by a system administrator.
Type	<p>The type of port channel:</p> <ul style="list-style-type: none"> Dynamic: Uses Link Aggregation Control Protocol (LACP) Protocol Data Units (PDUs) to exchange information with the link partners to help maintain the link state. To utilize Dynamic link aggregation on this port channel, the link partner must also support LACP. Static: Does not require a partner system to be able to aggregate its member ports. When a port is added to a port channel as a static member, it neither transmits nor receives LACP PDUs. <p>When configuring a port channel, use the Static Mode field to set the port channel type. If the Static Mode is disabled, the port channel type is Dynamic.</p>
Admin Mode	The administrative mode of the port channel. When disabled, the port channel does not send and receive traffic.
STP Mode	The spanning tree protocol (STP) mode of the port channel. When enabled, the port channel participates in the STP operation to help prevent network loops.
Link State	The current link status of the port channel, which can be Up, Up (SFP), or Down.
Link Trap	The link trap mode of the port channel. When enabled, a trap is sent to any configured SNMP receiver(s) when the link state of the port channel changes.

Item	Description
Members	The ports that are members of a port channel. Each port channel can have a maximum of 8 member ports. To add ports to the port channel, select one or more ports from the Port List field (CTRL + click to select multiple ports). Then, use the appropriate arrow icon to move the selected ports to the Members field.
Active Ports	The ports that are actively participating members of a port channel. A member port that is operationally or administratively disabled or does not have a link is not an active port.
Load Balance	The algorithm used to distribute traffic load among the physical ports of the port channel while preserving the per-flow packet order. The packet attributes the load-balancing algorithm can use to determine the outgoing physical port include the following: <ul style="list-style-type: none"> ■ Source MAC, VLAN, Ethertype, Incoming Port ■ Destination MAC, VLAN, Ethertype, Incoming Port ■ Source/Destination MAC, VLAN, Ethertype, Incoming Port ■ Source IP and Source TCP/UDP Port Fields ■ Destination IP and Destination TCP/UDP Port Fields ■ Source/Destination IP and TCP/UDP Port Fields
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.

4.4.17.2 Statistics

The Port Channel Statistics page displays the flap count for each port channel and their member ports. A flap occurs when a port-channel interface or port-channel member port goes down.

To access this page, click **Switching > Port Channel > Statistics**.

Interface	Channel Name	Type	Flap Count
1/1	ch1	Port Channel	0
1/2	ch2	Port Channel	0
1/3	ch3	Port Channel	0
1/4	ch4	Port Channel	0
1/5	ch5	Port Channel	0
1/6	ch6	Port Channel	0
1/7	ch7	Port Channel	0
1/8	ch8	Port Channel	0
1/9	ch9	Port Channel	0
1/10	ch10	Port Channel	0

Figure 4.264 Switching > Port Channel > Statistics

The following table describes the items in the previous figure.

Item	Description
Interface	The port channel or member port (physical port) associated with the rest of the data in the row.
Channel Name	The port channel name associated with the port channel. For a physical port, this field identifies the name of the port channel of which the port is a member.
Type	The interface type, which is either Port Channel (logical link-aggregation group) or Member Port (physical port).

Item	Description
Flap Count	The number of times the interface has gone down. The counter for a member port is incremented when the physical port is either manually shut down by the administrator or when its link state is down. When a port channel is administratively shut down, the flap counter for the port channel is incremented, but the flap counters for its member ports are not affected. When all active member ports for a port channel are inactive (either administratively down or link down), then the port channel flap counter is incremented.
Refresh	Click Refresh to update the screen.
Clear Counters	Click Clear Counters to reset the flap counters for all port channels and member ports to zero.

4.4.18 Port Security

Port Security can be enabled on a per-port basis. When a port is locked, only packets with allowable source MAC addresses can be forwarded. All other packets are discarded. A MAC address can be defined as allowable by one of two methods: dynamically or statically. Note that both methods are used concurrently when a port is locked.

Dynamic locking implements a “first arrival” mechanism for Port Security. You specify how many addresses can be learned on the locked port. If the limit has not been reached, a packet with an unknown source MAC address is learned and forwarded normally. Once the limit is reached, no more addresses are learned on the port. Any packets with source MAC addresses that were not already learned are discarded. Note that you can effectively disable dynamic locking by setting the number of allowable dynamic entries to zero.

Static locking allows you to specify a list of MAC addresses that are allowed on a port. The behavior of packets is the same as for dynamic locking: only packets with an allowable source MAC address can be forwarded.

4.4.18.1 Global

Use the Port Security Global Administration page to configure the global administrative mode for the port security feature. Port security, which is also known as port MAC locking, allows you to limit the number of source MAC address that can be learned on a port. If a port reaches the configured limit, any other addresses beyond that limit are not learned, and the frames are discarded. Frames with a source MAC address that has already been learned will be forwarded. Port security can help secure the network by preventing unknown devices from forwarding packets into the network.

To access this page, click **Switching > Port Security > Global**.

Figure 4.265 Switching > Port Security > Global

The following table describes the items in the previous figure.

Item	Description
Port Security Admin Mode	Enable or disable the global administrative mode for port security. The port security mode must be enabled both globally and on an interface to enforce the configured limits for the number of static and dynamic MAC addresses allowed on that interface.
Submit	Click Submit to save the values and update the screen.

Item	Description
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.4.18.2 Interface

Use the Port Security Interface Status page to configure the port security feature on a selected interface.

To access this page, click **Switching > Port Security > Interface**.

The screenshot shows a web interface for configuring port security. At the top, it says 'Display 10 rows' and 'Showing 1 to 10 of 42 entries'. There is a search filter box. Below is a table with columns: Interface, Port Security Mode, Max Dynamic Addresses Allowed, Max Static Addresses Allowed, Sticky Mode, Violation Trap Mode, Violation Shutdown Mode, and Last Violation MAC/VLAN. The table lists 10 interfaces (0/1 to 0/10) with 'Disable' for Port Security Mode, '600' for Max Dynamic Addresses Allowed, and '20' for Max Static Addresses Allowed. All other modes are set to 'Disable'. At the bottom, there are navigation buttons: 'First', 'Previous', '1', '2', '3', '4', '5', 'Next', 'Last', and action buttons: 'Refresh', 'Edit', and 'Edit All'.

Interface	Port Security Mode	Max Dynamic Addresses Allowed	Max Static Addresses Allowed	Sticky Mode	Violation Trap Mode	Violation Shutdown Mode	Last Violation MAC/VLAN
0/1	Disable	600	20	Disable	Disable	Disable	
0/2	Disable	600	20	Disable	Disable	Disable	
0/3	Disable	600	20	Disable	Disable	Disable	
0/4	Disable	600	20	Disable	Disable	Disable	
0/5	Disable	600	20	Disable	Disable	Disable	
0/6	Disable	600	20	Disable	Disable	Disable	
0/7	Disable	600	20	Disable	Disable	Disable	
0/8	Disable	600	20	Disable	Disable	Disable	
0/9	Disable	600	20	Disable	Disable	Disable	
0/10	Disable	600	20	Disable	Disable	Disable	

Figure 4.266 Switching > Port Security > Interface

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. When configuring the port security settings for one or more interfaces, this field lists the interfaces that are being configured.
Port Security Mode	The administrative mode of the port security feature on the interface. The port security mode must be enabled both globally and on an interface to enforce the configured limits for the number of static and dynamic MAC addresses allowed on that interface.
Max Dynamic Addresses Allowed	The number of source MAC addresses that can be dynamically learned on an interface. If an interface reaches the configured limit, any other addresses beyond that limit are not learned, and the frames are discarded. Frames with a source MAC address that has already been learned will be forwarded. A dynamically-learned MAC address is removed from the MAC address table if the entry ages out, the link goes down, or the system resets. Note that the behavior of a dynamically-learned address changes if the sticky mode for the interface is enabled or the address is converted to a static MAC address.
Max Static Addresses Allowed	The number of source MAC addresses that can be manually added to the port security MAC address table for an interface. If the port link goes down, the statically configured MAC addresses remain in the MAC address table. The maximum number includes all dynamically-learned MAC addresses that have been converted to static MAC addresses.

Item	Description
Sticky Mode	The sticky MAC address learning mode, which is one of the following: <ul style="list-style-type: none"> Enabled: MAC addresses learned or manually configured on this interface are learned in sticky mode. A sticky-mode MAC address is a MAC address that does not age out and is added to the running configuration. If the running configuration is saved to the startup configuration, the sticky addresses are saved to persistent storage and do not need to be relearned when the device restarts. Upon enabling sticky mode on an interface, all dynamically learned MAC addresses in the MAC address table for that interface are converted to sticky mode. Additionally, new addresses dynamically learned on the interface will also become sticky. Disabled: When a link goes down on a port, all of the dynamically learned addresses are cleared from the source MAC address table the feature maintains. When the link is restored, the interface can once again learn addresses up to the specified limit. If sticky mode is disabled after being enabled on an interface, the sticky-mode addresses learned or manually configured on the interface are converted to dynamic entries and are automatically removed from persistent storage.
Violation Trap Mode	Indicates whether the port security feature sends a trap to the SNMP agent when a port is locked and a frame with a MAC address not currently in the table arrives on the port. A port is considered to be locked once it has reached the maximum number of allowed dynamic or static MAC address entries in the port security MAC address table.
Violation Shutdown Mode	Indicates whether the port security feature shuts down the port after MAC limit is reached.
Last Violation MAC/VLAN	The source MAC address and, if applicable, associated VLAN ID of the last frame that was discarded at a locked port.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.
Edit All	Click Edit All to apply same settings to all interfaces.

4.4.18.3 VLAN

VLAN MAC Locking allows a network administrator to secure the network by locking down allowable MAC addresses on a given VLAN. Packets with a matching source MAC address can be forwarded normally. All other packets will be discarded.

To access this page, click **Switching > Port Security > VLAN**.




Figure 4.267 Switching > Port Security > VLAN

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The VLAN ID specified in the Ethernet frame received by the interface.
Interface	The interface associated with the rest of the data in the row.
Dynamic MAC Address	The MAC address that was learned on the device. An address is dynamically learned when a frame arrives on the interface and the source MAC address in the frame is added to the MAC address table.

Item	Description
Max Dynamic Addresses Allowed	The number of source MAC addresses that can be dynamically learned on an interface. If an interface reaches the configured limit, any other addresses beyond that limit are not learned, and the frames are discarded. Frames with a source MAC address that has already been learned will be forwarded. A dynamically-learned MAC address is removed from the MAC address table if the entry ages out, the link goes down, or the system resets. Note that the behavior of a dynamically-learned address changes if the sticky mode for the interface is enabled or the address is converted to a static MAC address.
Operational MAC Limit	The number of source MAC addresses that are dynamically currently reached to that of Maximum Configured MAC Limit.
Violation Shutdown Mode	After MAC limit has reached, action will shut down the ports participating in the VLAN.
Violation Trap Mode	After MAC limit has reached, a log message will be generated with violation MAC address details.
Refresh	Click Refresh to update the screen.
Add	Click Add to configure VLANs. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

To configure VLANs:

Click **Switching > Port Security > VLAN > Add**.




Figure 4.268 Switching > Port Security > VLAN > VLAN

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The VLAN ID specified in the Ethernet frame received by the interface.
Max Dynamic Addresses Allowed	The number of source MAC addresses that can be dynamically learned on an interface. If an interface reaches the configured limit, any other addresses beyond that limit are not learned, and the frames are discarded. Frames with a source MAC address that has already been learned will be forwarded. A dynamically-learned MAC address is removed from the MAC address table if the entry ages out, the link goes down, or the system resets. Note that the behavior of a dynamically-learned address changes if the sticky mode for the interface is enabled or the address is converted to a static MAC address.
Violation Shutdown Mode	After MAC limit has reached, action will shut down the ports participating in the VLAN.
Violation Trap Mode	After MAC limit has reached, a log message will be generated with violation MAC address details.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.4.18.4 Static MAC

Use the Port Security Static MAC Addresses page to add and remove the MAC addresses of hosts that are allowed to send traffic to specific interfaces on the device. The number of MAC addresses you can associate with each interface is determined by the maximum static MAC addresses allowed on a given interface.

To access this page, click **Switching > Port Security > Static MAC**.




Figure 4.269 Switching > Port Security > Static MAC

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. When adding a static MAC address entry, use the Interface menu to select the interface to associate with the permitted MAC address.
Static MAC Address	The MAC address of the host that is allowed to forward packets on the associated interface.
VLAN ID	The ID of the VLAN that includes the host with the specified MAC address.
Sticky Mode	Indicates whether the static MAC address entry is added in sticky mode. When adding a static MAC address entry, the Sticky Mode field can be selected only if it is enabled on the interface. If a static MAC address is added in sticky mode, and sticky mode is disabled on the interface, the MAC address entry is converted to a dynamic entry and will age out and be removed from the running (and saved) configuration if it is not relearned.
Refresh	Click Refresh to update the screen.
Add	Click Add to associate a static MAC address with an interface. See the following procedure.
Remove	Click Remove to remove the selected entries.

To associate a static MAC address with an interface:

Click **Switching > Port Security > Static MAC > Add**.




Figure 4.270 Switching > Port Security > Static MAC > Add

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. When adding a static MAC address entry, use the Interface menu to select the interface to associate with the permitted MAC address.
Static MAC Address	The MAC address of the host that is allowed to forward packets on the associated interface.

Item	Description
VLAN ID	The ID of the VLAN that includes the host with the specified MAC address.
Sticky Mode	Indicates whether the static MAC address entry is added in sticky mode. When adding a static MAC address entry, the Sticky Mode field can be selected only if it is enabled on the interface. If a static MAC address is added in sticky mode, and sticky mode is disabled on the interface, the MAC address entry is converted to a dynamic entry and will age out and be removed from the running (and saved) configuration if it is not relearned.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.4.18.5 Dynamic MAC

Use the Port Security Dynamic MAC Addresses page to add and remove the MAC addresses of hosts that are allowed to send traffic to specific interfaces on the device. The number of MAC addresses you can associate with each interface is determined by the maximum static MAC addresses allowed on a given interface.

To access this page, click **Switching > Port Security > Dynamic MAC**.



Figure 4.271 Switching > Port Security > Dynamic MAC

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. When converting dynamic addresses to static addresses, use the Interface menu to select the interface to associate with the MAC addresses.
Dynamic MAC Address	The MAC address that was learned on the device. An address is dynamically learned when a frame arrives on the interface and the source MAC address in the frame is added to the MAC address table.
VLAN ID	The VLAN ID specified in the Ethernet frame received by the interface.
Refresh	Click Refresh to update the screen.
Convert to Static	Click Convert to Static to convert all MAC addresses learned on an interface to static MAC address entries.

4.4.19 Protected Ports

4.4.19.1 Configuration

Use the Protected Ports Configuration page to configure and view protected ports groups. A port that is a member of a protected ports group is a protected port. A port that is not a member of any protected ports group is an unprotected port. Each port can be a member of only one protected ports group. Ports in the same protected ports group cannot forward traffic to other protected ports within the group, even if they are members of the same VLAN. However, a port in a protected ports group can forward traffic to ports that are in a different protected ports group. A protected port can also forward traffic to unprotected ports. Unprotected ports can forward traffic to both protected and unprotected ports.

To access this page, click **Switching > Protected Ports > Configuration**.



Figure 4.272 Switching > Protected Ports > Configuration

The following table describes the items in the previous figure.

Item	Description
Group Name	The user-configured name of the protected ports group.
Protected Ports	The ports that are members of the protected ports group. When adding a port to a protected ports group, the Available Interfaces field lists the ports that are not already members of a protected ports group. To move an interface between the Available Interfaces and Selected Interfaces fields, click the port (or CTRL + click to select multiple ports), and then click the appropriate arrow to move the port(s) to the desired field.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new protected ports group and add ports to the group. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

To add a new protected ports group and add ports to the group:

Click **Switching > Protected Ports > Configuration > Add**.

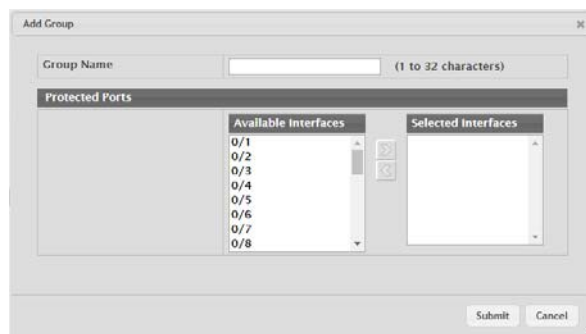


Figure 4.273 Switching > Protected Ports > Configuration > Add

The following table describes the items in the previous figure.

Item	Description
Group Name	The user-configured name of the protected ports group.

Item	Description
Protected Ports	The ports that are members of the protected ports group. When adding a port to a protected ports group, the Available Interfaces field lists the ports that are not already members of a protected ports group. To move an interface between the Available Interfaces and Selected Interfaces fields, click the port (or CTRL + click to select multiple ports), and then click the appropriate arrow to move the port(s) to the desired field.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.4.20 Spanning Tree

The Spanning Tree Protocol (STP) provides a tree topology for any arrangement of bridges. STP also provides one path between end stations on a network, eliminating loops. Spanning tree versions supported include Common STP, Multiple STP, and Rapid STP.

Classic STP provides a single path between end stations, avoiding and eliminating loops.

Multiple Spanning Tree Protocol (P) supports multiple instances of Spanning Tree to efficiently channel VLAN traffic over different interfaces. Each instance of the Spanning Tree behaves in the manner specified in IEEE 802.1w, Rapid Spanning Tree (RSTP), with slight modifications in the working but not the end effect (chief among the effects, is the rapid transitioning of the port to 'Forwarding'). The difference between the RSTP and the traditional STP (IEEE 802.1D) is the ability to configure and recognize full duplex connectivity and ports which are connected to end stations, resulting in rapid transitioning of the port to 'Forwarding' state and the suppression of Topology Change Notification. These features are represented by the parameters 'pointtopoint' and 'edgeport'. MSTP is compatible to both RSTP and STP. It behaves appropriately to STP and RSTP bridges. A MSTP bridge can be configured to behave entirely as a RSTP bridge or a STP bridge.

Note! *For two bridges to be in the same region, the force version should be 802.1S and their configuration name, digest key, and revision level should match. For more information about regions and their effect on network topology, refer to the IEEE 802.1Q standard.*



4.4.20.1 Switch

The Spanning Tree Switch Configuration page contains fields for enabling STP on the switch.

To access this page, click **Switching > Spanning Tree > Switch**.

Figure 4.274 Switching > Spanning Tree > Switch

The following table describes the items in the previous figure.

Item	Description
Spanning Tree Admin Mode	The administrative mode of STP on the device. When enabled, the device participates in the root bridge election process and exchanges Bridge Protocol Data Units (BPDUs) with other switches in the spanning tree to determine the root path costs and maintain topology information.
Force Protocol Version	The STP version the device uses, which is one of the following: <ul style="list-style-type: none"> IEEE 802.1d: Classic STP provides a single path between end stations, avoiding and eliminating loops. IEEE 802.1w: Rapid Spanning Tree Protocol (RSTP) behaves like classic STP but also has the ability to configure and recognize full-duplex connectivity and ports that are connected to end stations, resulting in rapid transitioning of the port to the Forwarding state and the suppression of Topology Change Notifications. IEEE 802.1s: Multiple Spanning Tree Protocol (MSTP) includes all the advantages of RSTP and also supports multiple spanning tree instances to efficiently channel VLAN traffic over different interfaces. MSTP is compatible with both RSTP and STP. PVST: Per-VLAN Spanning Tree (PVST) maintains a spanning tree instance for each VLAN configured in the network. This is based on IEEE 802.1d standard with additional features. RPVST: Rapid Per-VLAN Spanning Tree (RPVST) maintains a spanning tree instance for each VLAN configured in the network. This is based on IEEE 802.1w standard with additional features.
Configuration Name	The name of the MSTP region. Each switch that participates in the same MSTP region must share the same Configuration Name, Configuration Revision Level, and MST-to-VLAN mappings.
Configuration Revision Level	The revision number of the MSTP region. This number must be the same on all switches that participate in the MSTP region.
Configuration Digest Key	The 16 byte signature of type HMAC-MD5 created from the MST Configuration Table (a VLAN ID-to-MST ID mapping).
Configuration Format Selector	The version of the configuration format being used in the exchange of BPDUs.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.4.20.2 MST

Use the Spanning Tree MST Summary page to view and configure the Multiple Spanning Tree Instances (MSTIs) on the device. Multiple Spanning Tree Protocol (MSTP) allows the creation of MSTIs based upon a VLAN or groups of VLANs. Configuring MSTIs creates an active topology with a better distribution of network traffic and an increase in available bandwidth when compared to classic STP.

To access this page, click **Switching > Spanning Tree > MST**.




Figure 4.275 Switching > Spanning Tree > MST

The following table describes the items in the previous figure.

Item	Description
MST ID	The number that identifies the MST instance.
Priority	The bridge priority for the spanning-tree instance. This value affects the likelihood that the bridge is selected as the root bridge. A lower value increases the probability that the bridge is selected as the root bridge.
# of Associated VLANs	The number of VLANs that are mapped to the MSTI. This number does not contain any information about the VLAN IDs that are mapped to the instance.
Bridge Identifier	A unique value that is automatically generated based on the bridge priority value of the MSTI and the base MAC address of the bridge. When electing the root bridge for an MST instance, if the bridge priorities for multiple bridges are equal, the bridge with the lowest MAC address is elected as the root bridge.
Time Since Topology Change	The amount of time that has passed since the topology of the MSTI has changed.
Designated Root	The bridge identifier of the root bridge for the MST instance. The identifier is made up of the bridge priority and the base MAC address.
Root Path Cost	The path cost to the designated root for this MST instance. Traffic from a connected device to the root bridge takes the least-cost path to the bridge. If the value is 0, the cost is automatically calculated based on port speed.
Root Port	The port on the bridge with the least-cost path to the designated root for the MST instance.
Refresh	Click Refresh to update the screen.
Add	Click Add to configure a new MSTI. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

To add a new MST:

Click **Switching > Spanning Tree > MST > Add**.

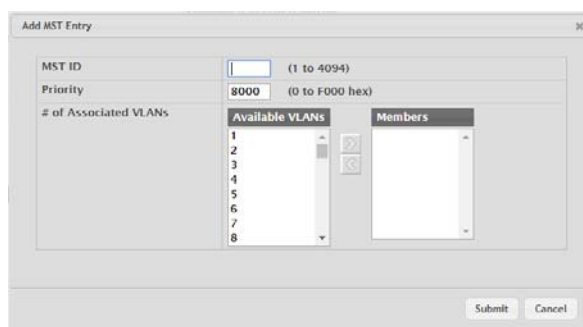


Figure 4.276 Switching > Spanning Tree > MST > Add

The following table describes the items in the previous figure.

Item	Description
MST ID	The number that identifies the MST instance.
Priority	The bridge priority for the spanning-tree instance. This value affects the likelihood that the bridge is selected as the root bridge. A lower value increases the probability that the bridge is selected as the root bridge.

Item	Description
# of Associated VLANs	The number of VLANs that are mapped to the MSTI. This number does not contain any information about the VLAN IDs that are mapped to the instance.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.4.20.3 MST Port

Use the Spanning Tree MST Port Summary page to view and configure the Multiple Spanning Tree (MST) settings for each interface on the device.

To access this page, click **Switching > Spanning Tree > MST Port**.

Interface	Port Role	Port Forwarding State	Port Priority	Port Path Cost	Description
ge0/1	Designated	Forwarding	0x0080	200000	
ge0/2	Disabled	Disabled	0x0080	0	
ge0/3	Disabled	Disabled	0x0080	0	
ge0/4	Disabled	Disabled	0x0080	0	
ge0/5	Disabled	Disabled	0x0080	0	
ge0/6	Disabled	Disabled	0x0080	0	
ge0/7	Disabled	Disabled	0x0080	0	
ge0/8	Disabled	Disabled	0x0080	0	
ge0/9	Disabled	Disabled	0x0080	0	
ge0/10	Disabled	Disabled	0x0080	0	

Figure 4.277 Switching > Spanning Tree > MST Port

The following table describes the items in the previous figure.

Item	Description
MST ID	The menu contains the ID of each MST instance that has been created on the device.
Interface	The port or link aggregation group (LAG) associated with the rest of the data in the row. When configuring MST settings for an interface, this field identifies the interface being configured.
Port Role	The role of the port within the MST, which is one of the following: <ul style="list-style-type: none"> ■ Root: A port on the non-root bridge that has the least-cost path to the root bridge. ■ Designated: A port that has the least-cost path to the root bridge on its segment. ■ Alternate: A blocked port that has an alternate path to the root bridge. ■ Backup: A blocked port that has a redundant path to the same network segment as another port on the bridge. ■ Master: The port on a bridge within an MST instance that links the MST instance to other STP regions. ■ Disabled: The port is administratively disabled and is not part of the spanning tree.

Item	Description
Port Forwarding State	<ul style="list-style-type: none"> ■ Blocking: The port discards user traffic and receives, but does not send, BPDUs. During the election process, all ports are in the blocking state. The port is blocked to prevent network loops. ■ Listening: The port sends and receives BPDUs and evaluates information to provide a loop-free topology. This state occurs during network convergence and is the first state in transitioning to the forwarding state. ■ Learning: The port learns the MAC addresses of frames it receives and begins to populate the MAC address table. This state occurs during network convergence and is the second state in transitioning to the forwarding state. ■ Forwarding: The port sends and receives user traffic. ■ Disabled: The port is administratively disabled and is not part of the spanning tree.
Port Priority	The priority for the port within the MSTI. This value is used in determining which port on a switch becomes the root port when two ports have the same least-cost path to the root. The port with the lower priority value becomes the root port. If the priority values are the same, the port with the lower interface index becomes the root port.
Port Path Cost	The path cost from the port to the root bridge.
Description	A user-configured description of the port.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.
Details	Click Details to open a window and display additional information for the selected interface.

4.4.20.4 CST

Use the Spanning Tree CST Configuration page to configure the Common Spanning Tree (CST) settings. The settings and information on this page define the device within the spanning tree topology that connects all STP/RSTP bridges and MSTP regions.

To access this page, click **Switching > Spanning Tree > CST**.

Bridge Priority	8000 (0 to F000 hex)
Bridge Max Age	20 (6 to 40)
Bridge Hello Time	2
Bridge Forward Delay	15 (4 to 30)
Spanning Tree Maximum Hops	20 (6 to 40)
BPDU Guard	<input type="checkbox"/>
BPDU Filter	<input type="checkbox"/>
Spanning Tree Tx Hold Count	6 (1 to 10)
Bridge Identifier	80:00:74:FE:48:20:BD:EC
Time Since Topology Change	0d:01:04:12
Topology Change Count	1
Topology Change	False
Designated Root	80:00:74:FE:48:20:BD:EC
Root Path Cost	0
Root Port	00:00
Max Age	20
Forward Delay	15
Hold Time	6
CST Regional Root	80:00:74:FE:48:20:BD:EC
CST Path Cost	0

Figure 4.278 Switching > Spanning Tree > CST

The following table describes the items in the previous figure.

Item	Description
Bridge Priority	The value that helps determine which bridge in the spanning tree is elected as the root bridge during STP convergence. A lower value increases the probability that the bridge becomes the root bridge.
Bridge Max Age	The amount of time a bridge waits before implementing a topological change.
Bridge Hello Time	The amount of time the root bridge waits between sending hello BPDUs.
Bridge Forward Delay	The amount of time a bridge remains in a listening and learning state before forwarding packets.
Spanning Tree Maximum Hops	The maximum number of hops a Bridge Protocol Data Unit (BPDU) is allowed to traverse within the spanning tree region before it is discarded.
BPDU Guard	When enabled, BPDU Guard can disable edge ports that receive BPDU packets. This prevents a new device from entering the existing STP topology. Thus devices that were originally not a part of STP are not allowed to influence the STP topology.
BPDU Filter	When enabled, this feature filters the BPDU traffic on the edge ports. When spanning tree is disabled on a port, BPDU filtering allows BPDU packets received on that port to be dropped.
Spanning Tree Tx Hold Count	The maximum number of BPDUs that a bridge is allowed to send within a hello time window.
Bridge Identifier	A unique value that is automatically generated based on the bridge priority value and the base MAC address of the bridge. When electing the root bridge for the spanning tree, if the bridge priorities for multiple bridges are equal, the bridge with the lowest MAC address is elected as the root bridge.
Time Since Topology Change	The amount of time that has passed since the topology of the spanning tree has changed since the device was last reset.
Topology Change Count	The number of times the topology of the spanning tree has changed.
Topology Change	Indicates whether a topology change is in progress on any port assigned to the CST. If a change is in progress the value is True; otherwise, it is False.
Designated Root	The bridge identifier of the root bridge for the CST. The identifier is made up of the bridge priority and the base MAC address.
Root Path Cost	The path cost to the designated root for the CST. Traffic from a connected device to the root bridge takes the least-cost path to the bridge. If the value is 0, the cost is automatically calculated based on port speed.
Root Port	The port on the bridge with the least-cost path to the designated root for the CST.
Max Age	The amount of time a bridge waits before implementing a topological change.
Forward Delay	The forward delay value for the root port bridge.
Hold Time	The minimum amount of time between transmissions of Configuration BPDUs.
CST Regional Root	The bridge identifier of the CST regional root. The identifier is made up of the priority value and the base MAC address of the regional root bridge.
CST Path Cost	The path cost to the CST tree regional root.
Submit	Click Submit to save the values and update the screen.

Item	Description
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.4.20.5 CST Port

Use the Spanning Tree CST Port Summary page to view and configure the Common Spanning Tree (CST) settings for each interface on the device.

To access this page, click **Switching > Spanning Tree > CST Port**.

Interface	Port Role	Port Forwarding State	Port Priority	Port Path Cost	Description
0/1	Disabled	Disabled	0x0080	0	
0/2	Disabled	Disabled	0x0080	0	
0/3	Disabled	Disabled	0x0080	0	
0/4	Disabled	Disabled	0x0080	0	
0/5	Disabled	Disabled	0x0080	0	
0/6	Disabled	Disabled	0x0080	0	
0/7	Disabled	Disabled	0x0080	0	
0/8	Disabled	Disabled	0x0080	0	
0/9	Disabled	Disabled	0x0080	0	
0/10	Designated	Forwarding	0x0080	0	

Figure 4.279 Switching > Spanning Tree > CST Port

The following table describes the items in the previous figure.

Item	Description
Interface	The port or link aggregation group (LAG) associated with the rest of the data in the row. When configuring CST settings for an interface, this field identifies the interface being configured.
Port Role	<p>The role of the port within the CST, which is one of the following:</p> <ul style="list-style-type: none"> ■ Root: A port on the non-root bridge that has the least-cost path to the root bridge. ■ Designated: A port that has the least-cost path to the root bridge on its segment. ■ Alternate: A blocked port that has an alternate path to the root bridge. ■ Backup: A blocked port that has a redundant path to the same network segment as another port on the bridge. ■ Master: The port on a bridge within an MST instance that links the MST instance to other STP regions. ■ Disabled: The port is administratively disabled and is not part of the spanning tree.
Port Forwarding State	<ul style="list-style-type: none"> ■ Blocking: The port discards user traffic and receives, but does not send, BPDUs. During the election process, all ports are in the blocking state. The port is blocked to prevent network loops. ■ Listening: The port sends and receives BPDUs and evaluates information to provide a loop-free topology. This state occurs during network convergence and is the first state in transitioning to the forwarding state. ■ Learning: The port learns the MAC addresses of frames it receives and begins to populate the MAC address table. This state occurs during network convergence and is the second state in transitioning to the forwarding state. ■ Forwarding: The port sends and receives user traffic. ■ Disabled: The port is administratively disabled and is not part of the spanning tree.

Item	Description
Port Priority	The priority for the port within the CST. This value is used in determining which port on a switch becomes the root port when two ports have the same least-cost path to the root. The port with the lower priority value becomes the root port. If the priority values are the same, the port with the lower interface index becomes the root port.
Port Path Cost	The path cost from the port to the root bridge.
Description	A user-configured description of the port.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.
Details	Click Details to open a window and display additional information for the selected interface.

4.4.20.6 Statistics

Use the Spanning Tree Statistics page to view information about the number and type of bridge protocol data units (BPDUs) transmitted and received on each port.

To access this page, click **Switching > Spanning Tree > Statistics**.

Interface	STP BPDUs Rx	STP BPDUs Tx	RSTP BPDUs Rx	RSTP BPDUs Tx	MSTP BPDUs Rx	MSTP BPDUs Tx	SSTP BPDUs Rx	SSTP BPDUs Tx
0/1	0	0	0	0	0	0	0	0
0/2	0	0	0	0	0	0	0	0
0/3	0	0	0	0	0	0	0	0
0/4	0	0	0	0	0	0	0	0
0/5	0	0	0	0	0	0	0	0
0/6	0	0	0	0	0	0	0	0
0/7	0	0	0	0	0	0	0	0
0/8	0	0	0	0	0	0	0	0
0/9	0	0	0	0	0	0	0	0
0/10	0	0	0	1041	0	0	0	0

Figure 4.280 Switching > Spanning Tree > Statistics

The following table describes the items in the previous figure.

Item	Description
Interface	The port or link aggregation group (LAG) associated with the rest of the data in the row.
STP BPDUs Rx	The number of classic STP (IEEE 802.1d) BPDUs received by the interface.
STP BPDUs Tx	The number of classic STP BPDUs sent by the interface.
RSTP BPDUs Rx	The number of RSTP (IEEE 802.1w) BPDUs received by the interface.
RSTP BPDUs Tx	The number of RSTP BPDUs sent by the interface.
MSTP BPDUs Rx	The number of MSTP (IEEE 802.1s) BPDUs received by the interface.
MSTP BPDUs Tx	The number of MSTP BPDUs sent by the interface.
Refresh	Click Refresh to update the screen.

4.4.20.7 PVST Global

Use the PVST/RPVST Global Configuration page to view and configure Per-VLAN Spanning Tree (PVST) / Rapid Per-VLAN Spanning Tree (RPVST) Global settings for the device.

To access this page, click **Switching > Spanning Tree > PVST Global**.

Figure 4.281 Switching > Spanning Tree > PVST Global

The following table describes the items in the previous figure.

Item	Description
Status	PVST configuration operational mode.
Fast Backbone	Configures Fast Backbone mode. When enabled, the switch detects the indirect link failures and accelerates the spanning tree convergence.
Fast Uplink	Configures Fast Uplink mode.
Max Update Rate (pps)	Configures Fast Uplink's Maximum Update Rate.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.4.20.8 PVST VLAN

Use the PVST/RPVST VLAN Configuration page to view and configure PVST/RPVST VLAN settings for the device.

To access this page, click **Switching > Spanning Tree > PVST VLAN**.

Figure 4.282 Switching > Spanning Tree > PVST VLAN

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The unique VLAN identifier (VID).
Root	Configures the switch to become the root bridge or standby root bridge by modifying the bridge priority to a lower value to ensure the bridge is the root (or standby) bridge.
Hello Time (Seconds)	The interval between sending successive BDPUs. Configures the spanning tree hello time interval for the specified VLAN.
Forward Delay (Seconds)	Configures the spanning tree forward delay time for a specified VLAN. This interval is the time spent in listening and learning states before transitioning a port to the forwarding states.
Max Age (Seconds)	The maximum age time before a bridge port saves its configuration information.
Bridge Priority	Configures the bridge priority of a VLAN. The value will be automatically adjusted (rounded) as needed by the system. If the value configured is not among the specified values, then it will be rounded off to the nearest valid value.

Item	Description
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new PVST VLAN. See the following procedure.
Deails	Click Details to open the PVST VLAN Details window.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

To add a new PVST VLAN:

Click **Switching > Spanning Tree > PVST VLAN > Add**.

Figure 4.283 Switching > Spanning Tree > PVST VLAN > Add

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The unique VLAN identifier (VID).
Root	Configures the switch to become the root bridge or standby root bridge by modifying the bridge priority to a lower value to ensure the bridge is the root (or standby) bridge.
Hello Time (Seconds)	The interval between sending successive BDPUs. Configures the spanning tree hello time interval for the specified VLAN.
Forward Delay (Seconds)	Configures the spanning tree forward delay time for a specified VLAN. This interval is the time spent in listening and learning states before transitioning a port to the forwarding states.
Max Age (Seconds)	The maximum age time before a bridge port saves its configuration information.
Bridge Priority	Configures the bridge priority of a VLAN. The value will be automatically adjusted (rounded) as needed by the system. If the value configured is not among the specified values, then it will be rounded off to the nearest valid value.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.4.20.9 PVST Interface

Use the PVST/RPVST Interface Configuration page to view and configure PVST/RPVST Interface settings for the device.

To access this page, click **Switching > Spanning Tree > PVST Interface**.

Figure 4.284 Switching > Spanning Tree > PVST Interface

The following table describes the items in the previous figure.

Item	Description
Interface	List of physical interfaces and LAGs.
Priority	The port priority configuration is used to allow the operator to select the relative importance of the port in the selection process for forwarding. Set this value to a lower number to prefer a port for forwarding of frames. This field is available for configuration only when PVST/RPVST is enabled.
Per VLAN Configuration	
VLAN ID	The unique VLAN identifier (VID).
Priority	The per VLAN priority value configuration of the port is the priority used to allow the operator to select the relative importance of the port in the selection process for forwarding. Set this value to a lower number to prefer a port for forwarding of frames. This priority configuration is used when the port is configured as a point-to-point link type.
Cost	The path cost from the port to the root bridge.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.

4.4.20.10 PVST Statistics

Use the PVST/RPVST Statistics page to view PVST/RPVST Statistics for the device. To access this page, click **Switching > Spanning Tree > PVST Statistics**.

Fast Backbone	
Transition via Fast Backbone	0
Inferior BPDUs Received	0
RLQ Request PDUs Received	0
RLQ Response PDUs Received	0
RLQ Request PDUs Sent	0
RLQ Response PDUs Sent	0
Fast Uplink	
Fast Uplink Transitions	0
Proxy Multicast Addresses Transmitted	0

Refresh

Figure 4.285 Switching > Spanning Tree > PVST Statistics

The following table describes the items in the previous figure.

Item	Description
Fast Backbone	
Transition via Fast Backbone	Number of fast backbone transitions.
Inferior BPDUs Received	Number of the received inferior BPDUs.
RLQ Request PDUs Received	Number of the received RLQ request PDUs.
RLQ Response PDUs Received	Number of the received RLQ response PDUs.
RLQ Request PDUs Sent	Number of the sent RLQ request PDUs.
RLQ Response PDUs Sent	Number of the sent RLQ response PDUs.
Fast Uplink	
Fast Uplink Transitions	Number of the fast uplink transitions.

Item	Description
Proxy Multicast Addresses Transmitted	Number of the transmitted proxy multicast addresses.
Refresh	Click Refresh to update the screen.

4.4.21 UDLD

4.4.21.1 Configuration

Use the UDLD Configuration page to configure the global Unidirectional Link Detection (UDLD) settings on the device. The UDLD feature detects unidirectional links on physical ports by exchanging packets containing information about neighboring devices. The purpose of the UDLD feature is to detect and avoid unidirectional links. A unidirectional link is a forwarding anomaly in a Layer 2 communication channel in which a bidirectional link stops passing traffic in one direction.

To access this page, click **Switching > UDLD > Configuration**.

The screenshot shows the UDLD Configuration page. At the top, there is a section for 'Admin Mode' with a radio button selected for 'Enable' and 'Disable' options. Below this, there are two input fields: 'Message Interval (Seconds)' with a value of '15' and a range of '(7 to 90)', and 'Timeout Interval (Seconds)' with a value of '5' and a range of '(5 to 60)'. At the bottom of the form, there are three buttons: 'Submit', 'Refresh', and 'Cancel'.

Figure 4.286 Switching > UDLD > Configuration

The following table describes the items in the previous figure.

Item	Description
Admin Mode	The administrative mode of UDLD on the device. UDLD must be administratively enabled on the device and on an interface for that interface to send UDLD messages. Additionally, UDLD must be enabled on the both sides of the link for the device to detect a unidirectional link.
Message Interval (Seconds)	The amount of time to wait between sending UDLD probe messages on ports that are in the advertisement phase.
Timeout Interval (Seconds)	The amount of time to wait to receive a UDLD message before considering the UDLD link to be unidirectional.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.4.21.2 Interface Configuration

Use the UDLD Interface Configuration page to configure the per-port UDLD settings. To access this page, click **Switching > UDLD > Interface Configuration**.

The screenshot shows a table with the following columns: 'Interface', 'Admin Mode', 'UDLD Mode', and 'UDLD Status'. The table contains four rows of data:

Interface	Admin Mode	UDLD Mode	UDLD Status
0/1	Disabled	Normal	Not Applicable
0/2	Disabled	Normal	Not Applicable
0/3	Disabled	Normal	Not Applicable
0/4	Disabled	Normal	Not Applicable

Figure 4.287 Switching > UDLD > Interface Configuration

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. In the Edit UDLD Interface Configuration window, this field identifies each interface that is being configured.

Item	Description
Admin Mode	The administrative mode of UDLD on the port.
UDLD Mode	<p>The UDLD mode for the port, which is one of the following:</p> <ul style="list-style-type: none"> ■ Normal: The state of the port is classified as Undetermined if an anomaly exists. An anomaly might be the absence of its own information in received UDLD messages or the failure to receive UDLD messages. An Undetermined state has no effect on the operation of the port. The port is not disabled and continues operating. When operating in UDLD normal mode, a port will be put into a disabled (Shutdown) state only in the following situations: <ul style="list-style-type: none"> – The UDLD PDU received from a partner does not have its own details (echo). – When there is a loopback, and information sent out on a port is received back exactly as it was sent. ■ Aggressive: The port is put into a disabled state for the same reasons that it occurs in normal mode. Additionally, a port in UDLD aggressive mode can be disabled if the port does not receive any UDLD echo packets even after bidirectional connection was established. If a bidirectional link is established, and packets suddenly stop coming from partner device, the UDLD aggressive-mode port assumes that link has become unidirectional.
UDLD Status	<p>The UDLD status on the port, which is one of the following:</p> <ul style="list-style-type: none"> ■ Not Applicable: The administrative status of UDLD is globally disabled or disabled on the interface. ■ Bidirectional: UDLD has detected a bidirectional link. ■ Shutdown: UDLD has detected a unidirectional link, and the port is in a disabled state. To clear the disabled state, click UDLD Port Reset. ■ Undetermined: UDLD has not collected enough information to determine the state of the port. ■ Unknown: The port link has physically gone down, but it is not because it was put in a disabled state by the UDLD feature.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to configure UDLD settings for one or more interfaces.
UDLD Port Reset	Click UDLD Port Reset to reset all UDLD ports that have a UDLD Status of Shutdown.

4.4.22 VLAN

Adding Virtual LAN (VLAN) support to a Layer 2 switch offers some of the benefits of both bridging and routing. Like a bridge, a VLAN switch forwards traffic based on the Layer 2 header, which is fast, and like a router, it partitions the network into logical segments, which provides better administration, security and management of multicast traffic.

A VLAN is a set of end stations and the switch ports that connect them. You may have many reasons for the logical division, such as department or project membership. The only physical requirement is that the end station and the port to which it is connected both belong to the same VLAN.

Each VLAN in a network has an associated VLAN ID, which appears in the IEEE 802.1Q tag in the Layer 2 header of packets transmitted on a VLAN. An end station may omit the tag, or the VLAN portion of the tag, in which case the first switch port to receive the packet may either reject it or insert a tag using its default VLAN ID. A given port may handle traffic for more than one VLAN, but it can only support one default VLAN ID.

4.4.22.1 Status

Use the VLAN Status page to view information about the VLANs configured on your system.

To access this page, click **Switching > VLAN > Status**.

Note! You cannot remove or rename VLAN 1.



VLAN ID	Name	Type	RSPAN
1	default	Default	

Figure 4.288 Switching > VLAN > Status

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The unique VLAN identifier (VID).
Name	A user-configurable name that identifies the VLAN.
Type	The type of VLAN, which can be one of the following: <ul style="list-style-type: none">■ Default: The default VLAN. This VLAN is always present, and the VLAN ID is 1.■ Static: A user-configured VLAN.■ Dynamic: A VLAN created by GARP VLAN Registration Protocol (GVRP).
RSPAN	Identifies whether the VLAN is configured (Enabled) as the Remote Switched Port Analyzer (RSPAN) VLAN. The RSPAN VLAN is used to carry mirrored traffic from source ports to a destination probe port on a remote device.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new VLAN. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

To add a new VLAN:

Click **Switching > VLAN > Status > Add**.




Figure 4.289 Switching > VLAN > Status > Add

The following table describes the items in the previous figure.

Item	Description
VLAN ID or Range	Specify VLAN ID(s). Use '-' to specify a range and ',' to separate VLAN IDs or VLAN ranges in the list.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.4.22.2 Port Configuration

Use the VLAN Port Configuration page to configure VLAN membership for the interfaces on the device and to specify whether traffic transmitted by the member ports should be tagged. The device supports IEEE 802.1Q tagging. Ethernet frames on a tagged VLAN have a 4-byte VLAN tag in the header.

To access this page, click **Switching > VLAN > Port Configuration**.

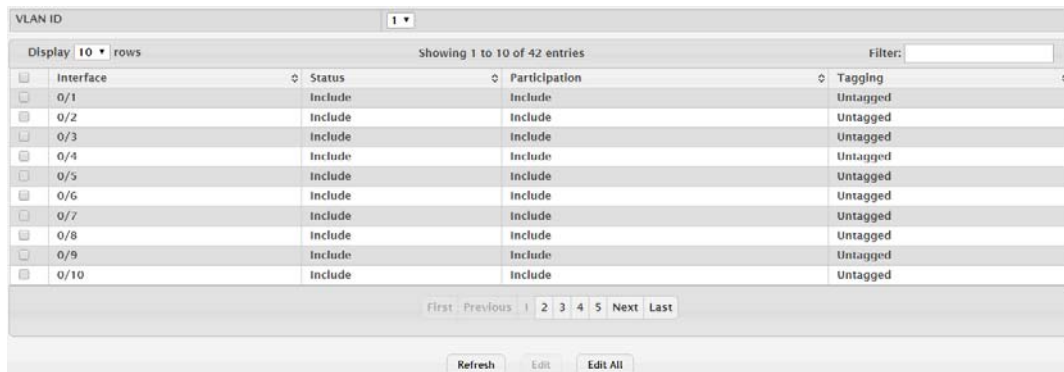


Figure 4.290 Switching > VLAN > Port Configuration

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The menu includes the VLAN ID for all VLANs configured on the device. Click the drop-down menu to select the correct VLAN to view or configure settings for a VLAN
Interface	The interface associated with the rest of the data in the row. When editing VLAN information for one or more interfaces, this field identifies the interfaces that are being configured.
Status	The current participation mode of the interface in the selected VLAN. The value of the Status field differs from the value of the Participation field only when the Participation mode is set to Auto Detect. The Status is one of the following: <ul style="list-style-type: none"> ■ Include: The port is a member of the selected VLAN. ■ Exclude: The port is not a member of the selected VLAN.

Item	Description
Participation	<p>The participation mode of the interface in the selected VLAN, which is one of the following:</p> <ul style="list-style-type: none"> ■ Include: The port is always a member of the selected VLAN. This mode is equivalent to registration fixed in the IEEE 802.1Q standard. ■ Exclude: The port is never a member of the selected VLAN. This mode is equivalent to registration forbidden in the IEEE 802.1Q standard. ■ Auto Detect: The port can be dynamically registered in the selected VLAN through GVRP or MVRP. The port will not participate in this VLAN unless it receives a GVRP or MVRP request and the device software supports the corresponding protocol. This mode is equivalent to registration normal in the IEEE 802.1Q standard.
Tagging	<p>The tagging behavior for all the ports in this VLAN, which is one of the following:</p> <ul style="list-style-type: none"> ■ Tagged: The frames transmitted in this VLAN will include a VLAN ID tag in the Ethernet header. ■ Untagged: The frames transmitted in this VLAN will be untagged.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected interfaces.
Edit All	Click Edit All to apply same settings to all interfaces.

4.4.22.3 Port Summary

Use the VLAN Port Summary page to configure the way interfaces handle VLAN-tagged, priority-tagged, and untagged traffic.

To access this page, click **Switching > VLAN > Port Summary**.

The screenshot shows a web interface for configuring VLAN settings. At the top, it says 'Display 10 rows' and 'Showing 1 to 10 of 42 entries'. There is a search filter box. Below is a table with columns: Interface, Port VLAN ID, Acceptable Frame Type, Ingress Filtering, Untagged VLANs, Tagged VLANs, Forbidden VLANs, Dynamic VLANs, and Priority. The table lists 10 interfaces (0/1 to 0/10) with 'Admit All' for Acceptable Frame Type, 'Disabled' for Ingress Filtering, and '1' for Untagged VLANs. At the bottom, there are navigation buttons: First, Previous, 1, 2, 3, 4, 5, Next, Last, and action buttons: Refresh, Edit, Edit All.

Interface	Port VLAN ID	Acceptable Frame Type	Ingress Filtering	Untagged VLANs	Tagged VLANs	Forbidden VLANs	Dynamic VLANs	Priority
0/1	1	Admit All	Disabled	1				0
0/2	1	Admit All	Disabled	1				0
0/3	1	Admit All	Disabled	1				0
0/4	1	Admit All	Disabled	1				0
0/5	1	Admit All	Disabled	1				0
0/6	1	Admit All	Disabled	1				0
0/7	1	Admit All	Disabled	1				0
0/8	1	Admit All	Disabled	1				0
0/9	1	Admit All	Disabled	1				0
0/10	1	Admit All	Disabled	1				0

Figure 4.291 Switching > VLAN > Port Summary

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. When editing information for one or more interfaces, this field identifies the interfaces that are being configured.
Port VLAN ID	The VLAN ID assigned to untagged or priority tagged frames received on this port. This value is also known as the Port VLAN ID (PVID). In a tagged frame, the VLAN is identified by the VLAN ID in the tag.

Item	Description
Acceptable Frame Type	<p>Indicates how the interface handles untagged and priority tagged frames. The options include the following:</p> <ul style="list-style-type: none"> ■ Admit All: Untagged and priority tagged frames received on the interface are accepted and assigned the value of the Port VLAN ID for this interface. ■ Only Tagged: The interface discards any untagged or priority tagged frames it receives. ■ Only Untagged: The interface discards any tagged frames it receives. <p>For all options, VLAN tagged frames are forwarded in accordance with the IEEE 802.1Q VLAN standard.</p>
Ingress Filtering	<p>Indicates how the interface handles tagged frames. The options include the following:</p> <ul style="list-style-type: none"> ■ Enabled: A tagged frame is discarded if this interface is not a member of the VLAN identified by the VLAN ID in the tag. ■ Disabled: All tagged frames are accepted.
Untagged VLANs	VLANs which are configured on the port to transmit egress packets as untagged.
Tagged VLANs	VLANs which are configured on the port to transmit egress packets as tagged.
Forbidden VLANs	When configuring port memberships in VLANs, you can specify one or more VLANs to be excluded from the available VLANs for the port. The forbidden VLANs list shows the VLANs to which the port cannot be assigned membership.
Dynamic VLANs	The list of VLANs of which the port became a member as result of the operations of dynamic VLAN protocols. When a VLAN is created as a dynamic VLAN, any port that is configured as switchport type Trunk or General automatically becomes a member of the VLAN, unless the VLAN port is excluded from the VLAN.
Priority	The default 802.1p priority assigned to untagged packets arriving at the interface.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected interfaces.
Edit All	Click Edit All to apply same settings to all interfaces.

4.4.22.4 Switchport Summary

Use the VLAN Switchport Summary page to configure switchport mode settings on interfaces. The switchport mode defines the purpose of the port based on the type of device it connects to and constraints the VLAN configuration of the port accordingly. Assigning the appropriate switchport mode helps simplify VLAN configuration and minimize errors.

To access this page, click **Switching > VLAN > Switchport Summary**.

The screenshot shows a web interface for configuring switchport modes. At the top, it says 'Display 10 rows' and 'Showing 1 to 10 of 42 entries'. Below this is a table with columns: Interface, Switchport Mode, Access VLAN ID, Native VLAN ID, Native VLAN Tagging, and Trunk Allowed VLANs. The table lists 10 interfaces (0/1 to 0/10), all in 'General' mode, with 'Access VLAN ID' and 'Native VLAN ID' set to 1, 'Native VLAN Tagging' set to 'Disabled', and 'Trunk Allowed VLANs' set to '1-4093'. At the bottom of the table are navigation buttons: 'First', 'Previous', '1', '2', '3', '4', '5', 'Next', 'Last', 'Refresh', 'Edit', and 'Edit All'.

Figure 4.292 Switching > VLAN > Switchport Summary

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. When editing information for one or more interfaces, this field identifies the interfaces that are being configured.
Switchport Mode	The switchport mode of the interface, which is one of the following: <ul style="list-style-type: none"> ■ Access: Access mode is suitable for ports connected to end stations or end users. Access ports participate only in one VLAN. They accept both tagged and untagged packets, but always transmit untagged packets. ■ Trunk: Trunk mode is intended for ports that are connected to other switches. Trunk ports can participate in multiple VLANs and accept both tagged and untagged packets. ■ General: General mode enables a custom configuration of a port. The user configures the General port VLAN attributes such as membership, PVID, tagging, ingress filter, etc., using the settings on the Port Configuration page. By default, all ports are initially configured in General mode.
Access VLAN ID	The access VLAN for the port, which is valid only when the port switchport mode is Access.
Native VLAN ID	The native VLAN for the port, which is valid only when the port switchport mode is Trunk.
Native VLAN Tagging	When enabled, if the trunk port receives untagged frames, it forwards them on the native VLAN with no VLAN tag. When disabled, if the port receives untagged frames, it includes the native VLAN ID in the VLAN tag when forwarding.
Trunk Allowed VLANs	The set of VLANs of which the port can be a member, when configured in Trunk mode. By default, this list contains all possible VLANs even if they have not yet been created.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected interfaces.
Edit All	Click Edit All to apply same settings to all interfaces.

4.4.22.5 Internal Usage

Use the VLAN Internal Usage page to configure which VLAN IDs to use for port-based routing interfaces. When a port-based routing interface is created, an unused VLAN ID is assigned internally. This page also displays a list of VLANs assigned to routing interfaces.

To access this page, click **Switching > VLAN > Internal Usage**.

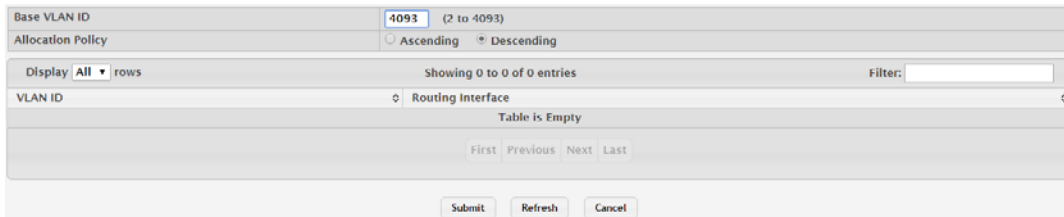


Figure 4.293 Switching > VLAN > Internal Usage

The following table describes the items in the previous figure.

Item	Description
Base VLAN ID	The first VLAN ID to be assigned to a port-based routing interface.
Allocation Policy	Determines whether VLAN IDs assigned to port-based routing interfaces start at the base and decrease in value (Descending) or start at the base and increase in value (Ascending).
VLAN ID	The VLAN ID assigned to a port-based routing interface. The device automatically assigns an unused VLAN ID when the routing interface is created.
Routing Interface	The port-based routing interface associated with the VLAN.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.4.22.6 Reset

Use the Reset VLAN Configuration page to reset all VLAN settings to their default values. Any VLANs that have been created on the system will be deleted.

To access this page, click **Switching > VLAN > Reset**.



Figure 4.294 Switching > VLAN > Reset

The following table describes the items in the previous figure.

Item	Description
Reset	Click Reset to initiates the action to reset all VLAN configuration parameters to their factory default settings.

4.4.22.7 RSPAN

Use the RSPAN Configuration page to configure the VLAN to use as the Remote Switched Port Analyzer (RSPAN) VLAN. RSPAN allows you to mirror traffic from multiple source ports (or from all ports that are members of a VLAN) from different network devices and send the mirrored traffic to a destination port (a probe port connected to a network analyzer) on a remote device. The mirrored traffic is tagged with the RSPAN VLAN ID and transmitted over trunk ports in the RSPAN VLAN.

To access this page, click **Switching > VLAN > RSPAN**.




Figure 4.295 Switching > VLAN > Status

The following table describes the items in the previous figure.

Item	Description
VLAN IDs	The VLANs configured on the system that are not currently enabled as Private VLANs. To enable a VLAN as a RSPAN VLAN, click the VLAN ID to select it (or CTRL + click to select multiple VLAN IDs). Then, click the appropriate arrow to move the selected VLAN or VLANs to the RSPAN VLAN IDs window.
RSPAN VLAN IDs	The VLANs that are enabled as RSPAN VLAN. To disable a VLAN as a RSPAN VLAN, click the VLAN ID to select it (or CTRL + click to select multiple VLAN IDs). Then, click the appropriate arrow to move the selected VLAN or VLANs to the VLAN IDs window.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.4.23 IP Subnet Based VLAN

4.4.23.1 Status

Use the IP Subnet Based VLAN Status page to add, edit, and remove IP subnet-based VLANs. IP subnet-based VLANs allow incoming untagged packets to be assigned to a VLAN based on the source IP address of the packet. All hosts in the same subnet are members of the same VLAN.

To access this page, click **Switching > IP Subnet Based VLAN > Status**.




Figure 4.296 Switching > IP Subnet Based VLAN > Status

The following table describes the items in the previous figure.

Item	Description
IP Address	The network address for the IP subnet. All incoming untagged packets that have a source IP address within the defined subnetwork are placed in the same VLAN.
Subnet Mask	The subnet mask that defines the network portion of the IP address.

Item	Description
VLAN ID	The VLAN ID of the IP subnet-based VLAN. If the source IP address of untagged traffic received on any port or LAG is within the associated IP subnet, the traffic is tagged with this VLAN ID.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new IP subnet-based VLAN. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

To add a new IP subnet-based VLAN:

Click **Switching > IP Subnet Based VLAN > Status > Add**.

Figure 4.297 Switching > IP Subnet Based VLAN > Status > Add

The following table describes the items in the previous figure.

Item	Description
IP Address	The network address for the IP subnet. All incoming untagged packets that have a source IP address within the defined subnetwork are placed in the same VLAN.
Subnet Mask	The subnet mask that defines the network portion of the IP address.
VLAN ID	The VLAN ID of the IP subnet-based VLAN. If the source IP address of untagged traffic received on any port or LAG is within the associated IP subnet, the traffic is tagged with this VLAN ID.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.4.24 MAC Based VLAN

4.4.24.1 Status

Use the MAC Based VLAN Status page to add, edit, or remove MAC-based VLANs. MAC-based VLANs allow incoming untagged packets to be assigned to a VLAN based on the source MAC address of the packet. This type of VLAN is useful when a host might not always connect to the network through the same port but needs to be on the same VLAN.

To access this page, click **Switching > MAC Based VLAN > Status**.

Figure 4.298 Switching > MAC Based VLAN > Status

The following table describes the items in the previous figure.

Item	Description
MAC Address	The source MAC address of the host. All untagged traffic that includes this address in the source MAC address field of the Ethernet frame is placed in the associated VLAN.
VLAN ID	The VLAN ID of the MAC-based VLAN. If an untagged frame received on any port or LAG matches the associated MAC address, it is tagged with this VLAN ID.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new MAC-based VLAN. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

To add a new MAC-based VLAN:

Click **Switching > MAC Based VLAN > Status > Add**.



Figure 4.299 Switching > MAC Based VLAN > Status > Add

The following table describes the items in the previous figure.

Item	Description
MAC Address	The source MAC address of the host. All untagged traffic that includes this address in the source MAC address field of the Ethernet frame is placed in the associated VLAN.
VLAN ID	The VLAN ID of the MAC-based VLAN. If an untagged frame received on any port or LAG matches the associated MAC address, it is tagged with this VLAN ID.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.4.25 Protocol Based VLAN

4.4.25.1 Status

Use the Protocol Based VLAN Status page to add and remove Protocol-based Virtual Local Area Networks (PBVLANS). In a PBVLAN, traffic is bridged through specified ports based on the protocol. PBVLANS allow you to define a packet filter that the device uses as the matching criteria to determine whether a particular packet belongs to a particular VLAN. PBVLANS are most often used in environments where network segments contain hosts running multiple protocols. PBVLANS can help optimize network traffic patterns because protocol-specific broadcast messages are sent only to hosts that use the protocols specified in the PBVLAN.

To access this page, click **Switching > Protocol Based VLAN > Status**.




Figure 4.300 Switching > Protocol Based VLAN > Status

The following table describes the items in the previous figure.

Item	Description
Group Name	The user-configured name that identifies the PBVLAN group.
VLAN	The VLAN ID associated with the PBVLAN. VLAN tagging for the PBVLAN works as follows: <ul style="list-style-type: none">■ If the frame received over a port is tagged, normal processing takes place.■ If the frame received over a port is untagged, the frame type is matched according to the protocol(s) assigned to the group on that port.<ul style="list-style-type: none">– If a match is found, the frame is assigned the VLAN ID specified for the group.– If a match is not found, the frame is assigned the port VID (PVID) as its VLAN ID.
Protocol	The protocol or protocols to use as the match criteria for an Ethernet frame. The protocol is included in the two-byte EtherType field of the frame. When adding a PBVLAN, you can specify the EtherType hex value or (for IP, ARP, and IPX) the protocol keyword.
Interface	The interfaces that are members of the PBVLAN group. An interface can be a member of multiple groups as long as the same protocol is not specified in more than one group that includes the interface. When adding a PBVLAN group, use the Available Interfaces and Group Interfaces fields to configure the interfaces that are members of the PBVLAN group. To move an interface between the Available Interfaces and Group Interfaces fields, click the interface (or CTRL + click to select multiple interfaces), and then click the appropriate arrow to move the selected interfaces to the desired field.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new protocol based VLAN. See the following procedure.
Remove	Click Remove to remove the selected entries.

To add a new protocol based VLAN:

Click **Switching > Protocol Based VLAN > Status > Add**.

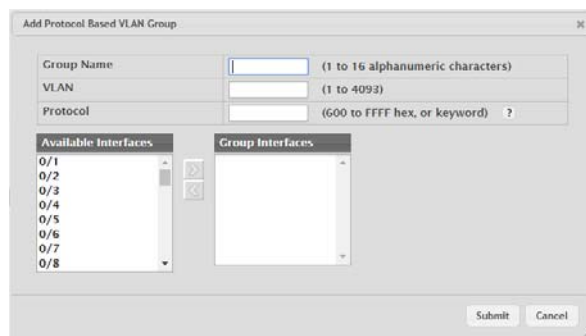


Figure 4.301 Switching > Protocol Based VLAN > Status > Add

The following table describes the items in the previous figure.

Item	Description
Group Name	The user-configured name that identifies the PBVLAN group.
VLAN	The VLAN ID associated with the PBVLAN. VLAN tagging for the PBVLAN works as follows: <ul style="list-style-type: none"> ■ If the frame received over a port is tagged, normal processing takes place. ■ If the frame received over a port is untagged, the frame type is matched according to the protocol(s) assigned to the group on that port. <ul style="list-style-type: none"> – If a match is found, the frame is assigned the VLAN ID specified for the group. – If a match is not found, the frame is assigned the port VID (PVID) as its VLAN ID.
Protocol	The protocol or protocols to use as the match criteria for an Ethernet frame. The protocol is included in the two-byte EtherType field of the frame. When adding a PBVLAN, you can specify the EtherType hex value or (for IP, ARP, and IPX) the protocol keyword.
Interface	The interfaces that are members of the PBVLAN group. An interface can be a member of multiple groups as long as the same protocol is not specified in more than one group that includes the interface. When adding a PBVLAN group, use the Available Interfaces and Group Interfaces fields to configure the interfaces that are members of the PBVLAN group. To move an interface between the Available Interfaces and Group Interfaces fields, click the interface (or CTRL + click to select multiple interfaces), and then click the appropriate arrow to move the selected interfaces to the desired field.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.4.25.2 Configuration

Use the Protocol Based VLAN Group Configuration page to configure existing Protocol-based VLAN (PBVLAN) groups. You can change the group name, VLAN ID, protocol information, and interfaces associated with the PBVLAN group.

To access this page, click **Switching > Protocol Based VLAN > Configuration**.

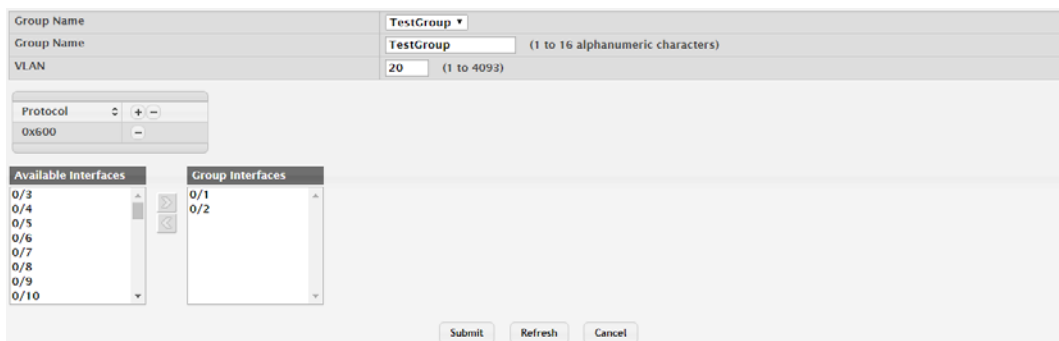


Figure 4.302 Switching > Protocol Based VLAN > Configuration

The following table describes the items in the previous figure.

Item	Description
Group Name	Click the drop-down menu select the PBLAN to change the properties.
Group Name	Enter the group name to update the name of the PBLAN group.
VLAN	The VLAN ID associated with the PBVLAN. Untagged traffic that matches the protocol criteria is tagged with this VLAN ID.
Protocol	<p>The protocol or protocols to use as the match criteria to determine whether a particular packet belongs to the PBVLAN. The protocols in this list are checked against the two-byte EtherType field of ingress Ethernet frames on the PVBLAN Group Interfaces. When adding a protocol, you can specify the EtherType hex value or (for IP, ARP, and IPX) the protocol keyword.</p> <p>To configure the protocols associated with a PBVLAN group, use the buttons available in the protocol table:</p> <ul style="list-style-type: none"> ■ To add a protocol to the group, click + button and enter the protocol to add. ■ To delete an entry from the list, click - button associated with the entry to remove. ■ To delete all entries from the list, click - button in the heading row.
Available Interfaces	The interfaces that can be added to the PBVLAN group. An interface can be a member of multiple groups as long as the same protocol is not specified in more than one group that includes the interface. To move an interface between the Available Interfaces and Group Interfaces fields, click the interface (or CTRL + click to select multiple interfaces), and then click the appropriate arrow to move the selected interfaces to the desired field.
Group Interfaces	The interfaces that are members of the PBVLAN group.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.4.26 Private VLAN

4.4.26.1 Configuration

Use the Private VLAN Configuration page to add Virtual Local Area Networks (VLANs) to the device and to configure existing VLANs as private VLANs. Private VLANs provide Layer 2 isolation between ports that share the same broadcast domain. In other words, a private VLAN allows a VLAN broadcast domain to be partitioned into smaller point-to-multipoint subdomains. The ports participating in a private VLAN can be located anywhere in the Layer 2 network. Each subdomain is defined (represented) by a primary VLAN and a secondary VLAN. The primary VLAN ID is the same for all subdomains that belong to a private VLAN. The secondary VLAN ID differentiates subdomains from each another and provides Layer 2 isolation between ports that are members of the same private VLAN.

Note! *The default VLAN and management VLAN are not displayed on the page because they cannot be configured as private VLANs.*



To access this page, click **Switching > Private VLAN > Configuration**.




Figure 4.303 Switching > Private VLAN > Configuration

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The ID of the VLAN that exists on the device.
Type	The private VLAN type, which is one of the following: <ul style="list-style-type: none">■ Unconfigured: The VLAN is not configured as a private VLAN.■ Primary: A private VLAN that forwards the traffic from the promiscuous ports to isolated ports, community ports, and other promiscuous ports in the same private VLAN. Only one primary VLAN can be configured per private VLAN. All ports within a private VLAN share the same primary VLAN.■ Isolated: A secondary VLAN that carries traffic from isolated ports to promiscuous ports. Only one isolated VLAN can be configured per private VLAN.■ Community: A secondary VLAN that forwards traffic between ports that belong to the same community and to the promiscuous ports. Multiple community VLANs can be configured per private VLAN.
Refresh	Click Refresh to update the screen.
Add VLAN	Click Add VLAN to add a new VLAN. See the following procedure.
Edit	Click Edit to edit the selected entries.

To add a new VLAN:

Click **Switching > Private VLAN > Configuration > Add VLAN**.




Figure 4.304 Switching > Private VLAN > Configuration > Add VLAN

The following table describes the items in the previous figure.

Item	Description
VLAN ID or Range	The ID of one or more VLANs to create. To create a single VLAN, enter its ID in the field. To create a continuous range of VLANs, use a hyphen (-) to separate the lowest and highest VLAN IDs in the range. To create multiple VLANs that are not in a continuous range, separate each VLAN ID or range of VLAN IDs with a comma (.). Do not use a space after the comma or anywhere in the field.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.4.26.2 Association

Use the Private VLAN Association page to configure the association between the primary VLAN and secondary VLANs. Associating a secondary VLAN with a primary VLAN allows host ports in the secondary VLAN to communicate outside the private VLAN.

Note! *Isolated VLANs and Community VLANs are collectively called Secondary VLANs.*



To access this page, click **Switching > Private VLAN > Association**.




Figure 4.305 Switching > Private VLAN > Association

The following table describes the items in the previous figure.

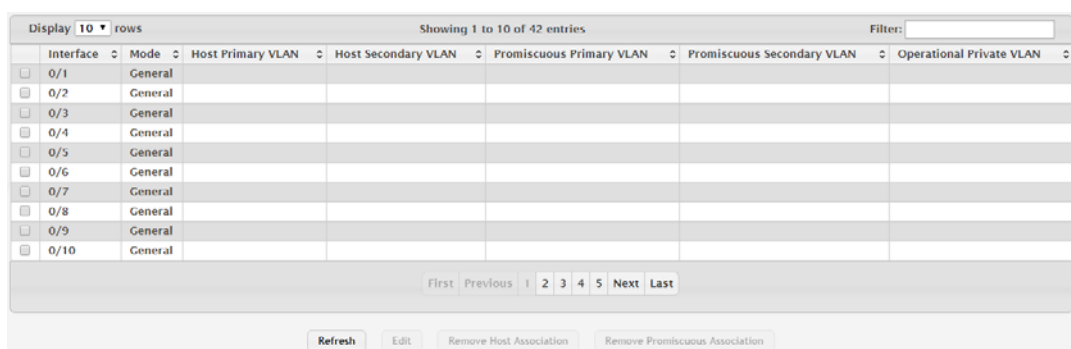
Item	Description
Primary VLAN	The VLAN ID of each VLAN configured as a primary VLAN.
Isolated VLAN	The VLAN ID of the isolated VLAN associated with the primary VLAN. If the field is blank, no isolated VLAN has been associated with the primary VLAN. An isolated VLAN is a secondary VLAN that carries traffic from isolated ports to promiscuous ports. Only one isolated VLAN can be configured per private VLAN.

Item	Description
Community VLAN	The VLAN ID of each community VLAN associated with the primary VLAN. If the field is blank, no community VLANs have been associated with the primary VLAN. A community VLAN is a secondary VLAN that forwards traffic between ports that belong to the same community and to the promiscuous ports. Multiple community VLANs can be configured per private VLAN.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.

4.4.26.3 Interface

Use the Private VLAN Interface Association page to configure the port mode for the ports and LAGs that belong to a private VLAN and to configure associations between interfaces and primary/secondary private VLANs.

To access this page, click **Switching > Private VLAN > Interface**.



Interface	Mode	Host Primary VLAN	Host Secondary VLAN	Promiscuous Primary VLAN	Promiscuous Secondary VLAN	Operational Private VLAN
0/1	General					
0/2	General					
0/3	General					
0/4	General					
0/5	General					
0/6	General					
0/7	General					
0/8	General					
0/9	General					
0/10	General					

Figure 4.306 Switching > Private VLAN > Interface

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. When editing interface settings, this field identifies the interface being configured.
Mode	The private VLAN mode of the interface, which is one of the following: <ul style="list-style-type: none"> ■ General: The interface is in general mode and is not a member of a private VLAN. ■ Promiscuous: The interface belongs to a primary VLAN and can communicate with all interfaces in the private VLAN, including other promiscuous ports, community ports, and isolated ports. ■ Host: The interface belongs to a secondary VLAN and, depending upon the type of secondary VLAN, can either communicate with other ports in the same community (if the secondary VLAN is a community VLAN) and with the promiscuous ports or is able to communicate only with the promiscuous ports (if the secondary VLAN is an isolated VLAN).
Host Primary VLAN	The primary private VLAN the port is a member of when it is configured to operate in Host mode.
Host Secondary VLAN	The secondary private VLAN the port is a member of when it is configured to operate in Host mode. The secondary private VLAN is either an isolated or community VLAN.
Promiscuous Primary VLAN	The primary private VLAN in which the port is a member when it is configured to operate in Promiscuous mode.

Item	Description
Promiscuous Secondary VLAN	The secondary private VLAN the port is a member of when it is configured to operate in Promiscuous mode. The secondary private VLAN is either an isolated or community VLAN.
Operational Private VLAN	The primary and secondary operational private VLANs for the interface. The VLANs that are operational depend on the configured mode for the interface and the private VLAN type.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.
Remove Host Association	Click Remove Host Association to remove the association between an interface and the primary/secondary private VLANs that the interface belongs to when it operates in host mode.
Remove Promiscuous Association	Click Remove Promiscuous Association to remove the association between an interface and the primary/secondary private VLANs that the interface belongs to when it operates in promiscuous mode.

4.4.27 Voice VLAN

4.4.27.1 Configuration

Use the Voice VLAN Configuration page to control the administrative mode of the Voice VLAN feature, which enables ports to carry voice traffic that has a defined priority. Voice over IP (VoIP) traffic is inherently time-sensitive: for a network to provide acceptable service, the transmission rate is vital. The priority level enables the separation of voice and data traffic entering the port.

To access this page, click **Switching > Voice VLAN > Configuration**.




Figure 4.307 Switching > Voice VLAN > Configuration

The following table describes the items in the previous figure.

Item	Description
Voice VLAN Admin Mode	The administrative mode of the Voice VLAN feature. When Voice VLAN is enabled globally and configured on interfaces that carry voice traffic, this feature can help ensure that the sound quality of an IP phone does not deteriorate when data traffic on the port is high.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.4.27.2 Interface Summary

Use this page to configure the per-port settings for the Voice VLAN feature. When Voice VLAN is configured on a port that receives both voice and data traffic, it can help ensure that the voice traffic has priority.

To access this page, click **Switching > Voice VLAN > Interface Summary**.




Figure 4.308 Switching > Voice VLAN > Interface Summary

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. When adding a Voice VLAN configuration to a port, the Interface menu allows you to select the port to configure. Only interfaces that have not been configured with Voice VLAN settings can be selected from the menu.
Operational State	The operational status of the Voice VLAN feature on the interface. To be enabled, Voice VLAN must be globally enabled and enabled on the interface. Additionally, the interface must be up and have a link.
CoS Override Mode	The Class of Service override mode: <ul style="list-style-type: none"> ■ Enabled: The port ignores the 802.1p priority value in the Ethernet frames it receives from connected devices. ■ Disabled: The port trusts the priority value in the received frame.
Voice VLAN Interface Mode	Indicates how an IP phone connected to the port should send voice traffic: <ul style="list-style-type: none"> ■ VLAN ID: Forward voice traffic in the specified voice VLAN. ■ Dot1p: Tag voice traffic with the specified 802.1p priority value. ■ None: Use the settings configured on the IP phone to send untagged voice traffic. ■ Untagged: Send untagged voice traffic. ■ Disable: Operationally disables the Voice VLAN feature on the interface.
Voice VLAN Interface Value	When adding or editing Voice VLAN settings for an interface and either VLAN ID or Dot1p is selected as the Voice VLAN Interface Mode, specify the voice VLAN ID or the Dot1p priority value that the connected IP phone should use for voice traffic.
Refresh	Click Refresh to update the screen.
Add	Click Add to configure Voice VLAN settings on a port. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

To configure Voice VLAN settings on a port:

Click **Switching > Voice VLAN > Interface Summary > Add**.




Figure 4.309 Switching > Voice VLAN > Interface Summary > Add

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. When adding a Voice VLAN configuration to a port, the Interface menu allows you to select the port to configure. Only interfaces that have not been configured with Voice VLAN settings can be selected from the menu.

Item	Description
CoS Override Mode	The Class of Service override mode: <ul style="list-style-type: none"> ■ Enabled: The port ignores the 802.1p priority value in the Ethernet frames it receives from connected devices. ■ Disabled: The port trusts the priority value in the received frame.
Voice VLAN Interface Mode	Indicates how an IP phone connected to the port should send voice traffic: <ul style="list-style-type: none"> ■ VLAN ID: Forward voice traffic in the specified voice VLAN. ■ Dot1p: Tag voice traffic with the specified 802.1p priority value. ■ None: Use the settings configured on the IP phone to send untagged voice traffic. ■ Untagged: Send untagged voice traffic. ■ Disable: Operationally disables the Voice VLAN feature on the interface.
Voice VLAN Interface Value	When adding or editing Voice VLAN settings for an interface and either VLAN ID or Dot1p is selected as the Voice VLAN Interface Mode, specify the voice VLAN ID or the Dot1p priority value that the connected IP phone should use for voice traffic.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.4.28 Virtual Port Channel

4.4.28.1 Global

Use the Virtual Port Channel Global Configuration page to view and manage global virtual port channel (VPC) settings on the device. VPCs are also known as multichassis or multiswitch link aggregation groups (MLAGs). Like port channels (also known as link aggregation groups or LAGs), VPCs allow one or more Ethernet links to be aggregated together to increase speed and provide redundancy. With port channels, the aggregated links must be on the same physical device, but VPCs do not share that requirement. The VPC feature allows links on two different switches to pair with links on a partner device. The partner device is unaware that it is pairing with two different devices to form a port channel.

To access this page, click **Switching > Virtual Port Channel > Global**.

Domain ID	<input type="text" value="1"/> (1 to 255)
VPC Mode	<input checked="" type="radio"/> Disable <input type="radio"/> Enable
Operational VPC Mode	Disabled
VPC State	Disabled
VPC MAC	<input type="text" value=""/> (xxxxxxxxxxxx)
Operational VPC MAC	
VPC System Priority	<input type="text" value="32767"/> (1 to 65535)
Operational VPC System Priority	
Self Role	None
Local System MAC	74FE48:20BD:EC
Keepalive Parameters	
Keepalive Priority	<input type="text" value="100"/> (1 to 255)
Keepalive Timeout (Seconds)	<input type="text" value="5"/> (2 to 15)
Keepalive Mode	<input checked="" type="radio"/> Disable <input type="radio"/> Enable
Peer	
Domain ID	
VPC MAC	
Operational VPC MAC	
VPC System Priority	
Operational VPC System Priority	
Peer Role	None
System MAC	
Peer Link	
Port Channel	<input type="text" value="None"/> <input checked="" type="checkbox"/> <input type="checkbox"/>
Peer Link Status	Down
Peer Link STP Mode	Disabled
Configured VLANs	None

Figure 4.310 Switching > Virtual Port Channel > Global

The following table describes the items in the previous figure.

Item	Description
Domain ID	The ID of the VPC domain. Only one VPC domain can be created on a given device. The VPC domain ID should be equal to the domain ID of the peer in order to form a VPC pair. The domain IDs are exchanged during role election and if different, VPC does not become operational.
VPC Mode	The administrative mode of VPC on the system.
Operational VPC Mode	The operational mode of VPC on the system. For VPC to be operational, several conditions must be met including the following: The VPC administrative mode is globally enabled. Peer links are configured. The Keepalive mode is enabled.
VPC State	The VPC state, which is one of the following: <ul style="list-style-type: none"> ■ Disable: The VPC mode is not operational. ■ Listen: The keepalive component does not advertise any packets. It listens for advertisements from a peer. ■ Ready: The keepalive component starts sending periodic keepalive messages. ■ Primary: Traffic over VPC interfaces is allowed to be forwarded in this state. The keepalive component continues to advertise keepalive messages with the state as Primary and monitors the health of the secondary device. ■ Secondary: Traffic over VPC interfaces is allowed to be forwarded in this state. The keepalive component continues to advertise keepalive messages with the state as Secondary and monitors the health of the primary device.
VPC MAC	The MAC address of the VPC domain. VPC MAC must be same on both the peer devices. The MAC address should be unicast and not be equal to the system MAC of either the primary or secondary VPC device. MAC addresses are exchanged during role election and if different, VPC does not become operational.
Operational VPC MAC	The VPC MAC address agreed upon by both peers during role election. This field is present in the keepalive message only if the transmitting peer is either primary or secondary.
VPC System Priority	The system priority of the VPC domain. System priority should be same on both the peer devices for VPC to become operational.
Operational VPC System Priority	The VPC system priority agreed upon by both peers during role election. This field is present in the keepalive message only if the transmitting peer is either primary or secondary.
Self Role	The role of the local device in the VPC domain, which is Primary, Secondary, or None. The role is determined by an election between the two devices after a keepalive link is established. The primary device owns the VPC member ports on the secondary device and handles the control plane functionality of supported protocols for the VPC member ports on the secondary device.
Local System MAC	The MAC address of the local system.
Keepalive Parameters	
Keepalive Priority	The priority value of the keepalive component on the local device. The device with lower priority value becomes the Primary device in the VPC role election.
Keepalive Timeout (Seconds)	The number of seconds that must pass without receiving a keepalive message before the peer device is considered to be down.
Keepalive Mode	The administrative mode of the keepalive component on the device.

Item	Description
Peer	
Domain ID	The ID of the peer VPC domain.
VPC MAC	The MAC address of the peer VPC domain.
Operational VPC MAC	The VPC MAC address agreed upon by both peers during role election.
VPC System Priority	The system priority of the peer VPC domain.
Operational VPC System Priority	The VPC system priority agreed upon by both peers during role election.
Peer Role	The role of the peer device in the VPC domain, which is Primary, Secondary, or None.
System MAC	The MAC address of the peer system.
Peer Link	
Port Channel	The port channel on the local device used for the peer link. To configure the peer link, click the Edit icon next to the field. The Edit window opens and allows you to select an available port channel from the Port Channel menu. To reset the port channel to the default value, click the Reset icon. The port channel cannot be changed or reset when the Operational VPC Mode is Enabled.
Peer Link Status	The operational status of the peer link, which is either Up or Down.
Peer Link STP Mode	The spanning tree protocol (STP) mode of the port channel. When enabled, the port channel participates in the STP operation to help prevent network loops.
Configured VLANs	The VLAN ID of each VLAN in which the port channel participates.
Egress Tagging	The VLAN ID tags included in the frames transmitted from the port channel.
Peer Detection	
Peer Detected	Indicates whether a peer link has been detected by DCPDP.
Peer IP Address	The IP address of the peer VPC device. This is the destination IP address in the DCPDP messages.
Source IP Address	The source IP address to be used by DCPDP.
UDP Port	The local UDP port to be used for listening to DCPDP packets.
Peer Detection Mode	The administrative mode of the peer detection feature (DCPDP).
Tx Interval	The interval in milliseconds between the DCPDP messages transmitted.
Operational Tx Interval	The operational transmit interval in milliseconds.
Rx Timeout	The DCPDP reception timeout in milliseconds.
Operational Rx Timeout	The operational timeout value in milliseconds.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Remove	Click Remove to remove the selected entries.
Cancel	Click Cancel to restore default value.

4.4.28.2 Interface Configuration

Use the Virtual Port Channel Interface Configuration page to configure the VPC interfaces on the device. A VPC interface is created by combining a port channel on the local device with a port channel on the peer device. The VPC interface on the local and peer devices share a common VPC identifier. You can configure multiple instances of VPC interfaces on each peer device in the VPC domain.

To access this page, click **Switching > Virtual Port Channel > Interface Configuration**.




Figure 4.311 Switching > Virtual Port Channel > Interface Configuration

The following table describes the items in the previous figure.

Item	Description
Interface	The ID of the local port channel configured as a VPC interface.
VPC Identifier	The VPC identifier of the interface. To form a VPC interface with a port channel on the peer device, the port channel on the peer device must use the same VPC identifier.
Operational Mode	The operational mode of the VPC interface.
VPC Interface State	The VPC interface state, which is one of the following: <ul style="list-style-type: none"> ■ Disabled: VPC functionality is operationally disabled on the VPC interface. ■ Wait: The port channel is waiting for VPC functionality to be enabled on a port channel on the peer device. ■ Error: VPC functionality is enabled on a port channel on both peer devices, but not all entry criteria are met for the port channel to be operational. For example, if the combined number of member ports for the VPC interface is more than the maximum allowed, then the state is set to Error on both devices. ■ Active: VPC functionality is enabled on a port channel on both peer devices, and all entry criteria are satisfied. The VPC interface is operationally enabled, and traffic is allowed to flow through the VPC member ports. ■ Inactive: The links connected to the VPC member ports are down, but the VPC interface on the peer remains active.
Configured VLANs	The VLAN ID of each VLAN in which the port channel participates.
Refresh	Click Refresh to update the screen.
Add	Click Add to configure a port channel as a VPC interface. See the following procedure.
Details	Click Details to open the Virtual Port Channel Interface Details window.
Remove	Click Remove to remove the selected entries.

To configure a port channel as a VPC interface:

Click **Switching > Virtual Port Channel > Interface Configuration > Add**.




Figure 4.312 Switching > Virtual Port Channel > Interface Configuration > Add
The following table describes the items in the previous figure.

Item	Description
Interface	The ID of the local port channel configured as a VPC interface.
VPC Identifier	The VPC identifier of the interface. To form a VPC interface with a port channel on the peer device, the port channel on the peer device must use the same VPC identifier.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.4.28.3 Statistics

The Virtual Port Channel Statistics page shows information about the number of messages of various types sent between the two VPC peer devices over the peer link.

To access this page, click **Switching > Virtual Port Channel > Statistics**.

Keepalive	
Total Transmitted	0
Tx Successful	0
Tx Errors	0
Total Received	0
Rx Successful	0
Rx Errors	0
Timeout Counter	0
Peer Link	
Control Messages Transmitted	0
Control Messages Tx Errors	0
Control Messages Tx Timeout (Seconds)	0
Control Messages ACK Transmitted	0
Control Messages ACK Tx Errors	0
Control Messages Received	0
Data Messages Transmitted	0
Data Messages Tx Errors	0
Data Messages Tx Timeout (Seconds)	0
Data Messages Received	0
BPDUs Transmitted To Peer	0
BPDUs Tx Error	0
BPDUs Received From Peer	0
BPDUs Rx Error	0
LACPDUs Transmitted To Peer	0
LACPDUs Tx Error	0
LACPDUs Received From Peer	0
LACPDUs Rx Error	0
<input type="button" value="Refresh"/> <input type="button" value="Clear Keepalive Counters"/> <input type="button" value="Clear Peer Link Counters"/>	

Figure 4.313 Switching > Virtual Port Channel > Statistics

The following table describes the items in the previous figure.

Item	Description
Keepalive	
Total Transmitted	The total number of keepalive messages the local device has sent to the peer device.
Tx Successful	The number of keepalive messages that have been successfully transmitted from the local device.

Item	Description
Tx Errors	The number of keepalive messages that the local device attempted to send to the peer device that were not transmitted due to an error.
Total Received	The total number of keepalive messages the local device has received from the peer device.
Rx Successful	The number of keepalive messages the local device has successfully received from the peer device.
Rx Errors	The number of keepalive messages the local device has received from the peer device that contained errors.
Timeout Counter	The number of times the keepalive timeout timer has expired.
Peer Link	
Control Messages Transmitted	The number of control messages successfully sent from the local device to the peer device over the peer link.
Control Messages Tx Errors	The number of errors encountered when sending peer-link control messages from the local device to the peer device over the peer link.
Control Messages Tx Timeout (Seconds)	The number of peer-link control messages that did not receive an ACK from the peer device.
Control Messages ACK Transmitted	The number of ACKs sent to the peer device in response to peer-link control messages that were received.
Control Messages ACK Tx Errors	The number of errors encountered when sending ACKs in response to peer-link control messages.
Control Messages Received	The number of control messages successfully received by the local device from the peer device over the peer link.
Data Messages Transmitted	The number of data messages successfully sent from the local device to the peer device over the peer link.
Data Messages Tx Errors	The number of errors encountered when sending peer-link data messages from the local device to the peer device over the peer link.
Data Messages Tx Timeout (Seconds)	The number of peer-link data messages that did not receive an ACK from the peer device.
Data Messages Received	The number of data messages successfully received by the local device from the peer device over the peer link.
BPDUs Transmitted To Peer	The number of BPDUs successfully sent to the peer device over the peer link.
BPDUs Tx Error	The number of errors encountered when sending BPDUs to the peer device.
BPDUs Received From Peer	The number of BPDUs successfully received from the peer device over the peer link.
BPDUs Rx Error	The number of errors encountered when receiving BPDUs from the peer device.
LACPDU Transmitted To Peer	The number of LACPDU s successfully sent to the peer device over the peer link.
LACPDU Tx Error	The number of errors encountered when sending LACPDU s to the peer device.
LACPDU Received From Peer	The number of LACPDU s successfully received from the peer device over the peer link.
LACPDU Received From Peer	The number of errors encountered when receiving LACPDU s from the peer device.
Refresh	Click Refresh to update the screen.
Clear Keepalive Counters	Click Clear Keepalive Counters to reset all keepalive message counters to 0.
Clear Peer Link Counters	Click Clear Peer Link Counters to reset all peer link message counters to 0.

4.5 Routing

When a packet enters the switch, the destination MAC address is checked to see if it matches any of the configured routing interfaces. If it does, then the silicon searches the host table for a matching destination IP address. If an entry is found, then the packet is routed to the host. If there is not a matching entry, then the switch performs a longest prefix match on the destination IP address. If an entry is found, then the packet is routed to the next hop. If there is no match, then the packet is routed to the next hop specified in the default route. If there is no default route configured, then the packet is passed to the 6200 series software to be handled appropriately.

The routing table can have entries added either statically by the administrator or dynamically via a routing protocol. The host table can have entries added either statically by the administrator or dynamically via ARP.

4.5.1 ARP Table

The ARP protocol associates a layer 2 MAC address with a layer 3 IPv4 address. FASTPATH SMB software features both dynamic and manual ARP configuration. With manual ARP configuration, you can statically add entries into the ARP table.

ARP is a necessary part of the internet protocol (IP) and is used to translate an IP address to a media (MAC) address, defined by a local area network (LAN) such as Ethernet. A station needing to send an IP packet must learn the MAC address of the IP destination, or of the next hop router, if the destination is not on the same subnet. This is achieved by broadcasting an ARP request packet, to which the intended recipient responds by unicasting an ARP reply containing its MAC address. Once learned, the MAC address is used in the destination address field of the layer 2 header prepended to the IP packet.

The ARP cache is a table maintained locally in each station on a network. ARP cache entries are learned by examining the source information in the ARP packet payload fields, regardless of whether it is an ARP request or response. Thus, when an ARP request is broadcast to all stations on a LAN segment or virtual LAN (VLAN), every recipient has the opportunity to store the sender's IP and MAC address in their respective ARP cache. The ARP response, being unicast, is normally seen only by the requestor, who stores the sender information in its ARP cache. Newer information always replaces existing content in the ARP cache.

The number of supported ARP entries is platform-dependent.

Devices can be moved in a network, which means the IP address that was at one time associated with a certain MAC address is now found using a different MAC, or may have disappeared from the network altogether (i.e., it has been reconfigured, disconnected, or powered off). This leads to stale information in the ARP cache unless entries are updated in reaction to new information seen on the network, periodically refreshed to determine if an address still exists, or removed from the cache if the entry has not been identified as a sender of an ARP packet during the course of an ageout interval, usually specified via configuration.

4.5.1.1 Summary

Use the ARP Table page to add an entry to the Address Resolution Protocol table.

To access this page, click **Routing > ARP Table > Summary**.




Figure 4.314 Routing > ARP Table > Summary

The following table describes the items in the previous figure.

Item	Description
IP Address	The IP address of a network host on a subnet attached to one of the device's routing interfaces. When adding a static ARP entry, specify the IP address for the entry after you click Add .
MAC Address	The unicast MAC address (hardware address) associated with the network host. When adding a static ARP entry, specify the MAC address to associate with the IP address in the entry.
Interface	The routing interface associated with the ARP entry. The network host is associated with the device through this interface.
Type	The ARP entry type: <ul style="list-style-type: none"> ■ Dynamic: An ARP entry that has been learned by the router ■ Gateway: A dynamic ARP entry that has the IP address of a routing interface ■ Local: An ARP entry associated with the MAC address of a routing interface on the device ■ Static: An ARP entry configured by the user
Age	The age of the entry since it was last learned or refreshed. This value is specified for Dynamic or Gateway entries only (it is left blank for all other entry types).
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new static ARP entry. See the following procedure.
Remove	Click Remove to remove the selected entries.

To add a new static ARP entry:

Click **Routing > ARP Table > Summary > Add**.



Figure 4.315 Routing > ARP Table > Summary > Add

The following table describes the items in the previous figure.

Item	Description
IP Address	The IP address of a network host on a subnet attached to one of the device's routing interfaces.
MAC Address	The unicast MAC address (hardware address) associated with the network host.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.5.1.2 Configuration

Use the ARP Table Configuration page to change the configuration parameters for the Address Resolution Protocol Table. You can also use this screen to display the contents of the table.

To access this page, click **Routing > ARP Table > Configuration**.

Age Time (Seconds)	1200	(15 to 21600)
Response Time (Seconds)	1	(1 to 10)
Retries	4	(0 to 10)
Cache Size	1536	(256 to 1536)
Dynamic Renew	<input type="checkbox"/>	

Figure 4.316 Routing > ARP Table > Configuration

The following table describes the items in the previous figure.

Item	Description
Age Time (Seconds)	The amount of time, in seconds, that a dynamic ARP entry remains in the ARP table before aging out.
Response Time (Seconds)	The amount of time, in seconds, that the device waits for an ARP response to an ARP request that it sends.
Retries	The maximum number of times an ARP request will be retried after an ARP response is not received. The number includes the initial ARP request.
Cache Size	The maximum number of entries allowed in the ARP table. This number includes all static and dynamic ARP entries.
Dynamic Renew	When selected, this option allows the ARP component to automatically attempt to renew dynamic ARP entries when they age out.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.5.1.3 Statistics

Use the ARP Table Statistics page to view the statistics for the Address Resolution Protocol Table. You can also use this screen to display the contents of the table.

To access this page, click **Routing > ARP Table > Statistics**.

Total Entry Count	0
Peak Total Entries	0
Active Static Entries	0
Configured Static Entries	0
Maximum Static Entries	64

Figure 4.317 Routing > ARP Table > Statistics

The following table describes the items in the previous figure.

Item	Description
Total Entry Count	The total number of entries currently in the ARP table. The number includes both dynamically learned entries and statically configured entries.
Peak Total Entries	The highest value reached by the Total Entry Count. This value is reset whenever the ARP table Cache Size configuration parameter is changed.

Item	Description
Active Static Entries	The total number of active ARP entries in the ARP table that were statically configured. After a static ARP entry is configured, it might not become active until certain other routing configuration conditions are met.
Configured Static Entries	The total number of static ARP entries that are currently in the ARP table. This number includes static ARP entries that are not active.
Maximum Static Entries	The maximum number of static ARP entries that can be configured in the ARP table.
Refresh	Click Refresh to update the screen.

4.5.2 IP

4.5.2.1 Configuration



Use the Routing IP Configuration page to configure global routing settings on the device. Routing provides a means of transmitting IP packets between subnets on the network. Routing configuration is necessary only if the device is used as a Layer 3 device that routes packets between subnets. If the device is used as a Layer 2 device that handles switching only, it typically connects to an external Layer 3 device that handles the routing functions; therefore, routing configuration is not required on the Layer 2 device.

To access this page, click **Routing > IP > Configuration**.

Figure 4.318 Routing > IP > Configuration

The following table describes the items in the previous figure.

Item	Description
Routing Mode	The administrative mode of routing on the device. The options are as follows: <ul style="list-style-type: none"> ■ Enable: The device can act as a Layer 3 device by routing packets between interfaces configured for IP routing. ■ Disable: The device acts as a Layer 2 bridge and switches traffic between interfaces. The device does not perform any internetwork routing.
ICMP Echo Replies	Select this option to allow the device to send ICMP Echo Reply messages in response to ICMP Echo Request (ping) messages it receives.
ICMP Redirects	Select this option to allow the device to send ICMP Redirect messages to hosts. An ICMP Redirect message notifies a host when a better route to a particular destination is available on the network segment.
ICMP Rate Limit Interval	The maximum burst interval for ICMP error messages transmitted by the device. The rate limit for ICMP error messages is configured as a token bucket. The ICMP Rate Limit Interval specifies how often the token bucket is initialized with tokens of the size configured in the ICMP Rate Limit Burst Size field.

Item	Description
ICMP Rate Limit Burst Size	The number of ICMP error messages that can be sent during the burst interval configured in the ICMP Rate Limit Interval field.
Static Route Preference	The default distance (preference) for static routes. Lower route-distance values are preferred when determining the best route. The value configured for Static Route Preference is used when using the CLI to configure a static route and no preference is specified. Changing the Static Route Preference does not update the preference of existing static routes.
Local Route Preference	The default distance (preference) for local routes.
Maximum Next Hops	The maximum number of hops the device supports.
Maximum Routes	The maximum number of routes that can exist in the routing table.
Global Default Gateway	The IP address of the default gateway for the device. If the destination IP address in a packet does not match any routes in the routing table, the packet is sent to the default gateway. The gateway specified in this field is more preferred than a default gateway learned from a DHCP server. Use the icons associated with this field to perform the following tasks: <ul style="list-style-type: none"> ■ To configure the default gateway, click  button and specify the IP address of the default gateway in the available field. ■ To reset the IP address of the default gateway to the factory default value, click  button associated with this field.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.5.2.2 Interface Summary

The Routing IP Interface Summary page shows summary information about the routing configuration for all interfaces.

To access this page, click **Routing > IP > Interface Summary**.



The screenshot shows a table with the following columns: Interface, Status, IP Address, Subnet Mask, Admin Mode, State, MAC Address, and IP MTU. The table lists 10 interfaces (ge0/1 to ge0/10). The status of all interfaces is 'Down'. The IP Address and Subnet Mask for all are '0.0.0.0'. The Admin Mode is 'Enabled' for all. The State is 'Active' for ge0/1 and 'Inactive' for the others. The MAC Address is '00:11:22:33:44:55' for all. The IP MTU is '1500' for all. Below the table are navigation buttons: First, Previous, 1, 2, Next, Last. At the bottom are buttons for Refresh, Edit, and Details.

Interface	Status	IP Address	Subnet Mask	Admin Mode	State	MAC Address	IP MTU
ge0/1	Down	0.0.0.0	0.0.0.0	Enabled	Active	00:11:22:33:44:55	1500
ge0/2	Down	0.0.0.0	0.0.0.0	Enabled	Inactive	00:11:22:33:44:55	1500
ge0/3	Down	0.0.0.0	0.0.0.0	Enabled	Inactive	00:11:22:33:44:55	1500
ge0/4	Down	0.0.0.0	0.0.0.0	Enabled	Inactive	00:11:22:33:44:55	1500
ge0/5	Down	0.0.0.0	0.0.0.0	Enabled	Inactive	00:11:22:33:44:55	1500
ge0/6	Down	0.0.0.0	0.0.0.0	Enabled	Inactive	00:11:22:33:44:55	1500
ge0/7	Down	0.0.0.0	0.0.0.0	Enabled	Inactive	00:11:22:33:44:55	1500
ge0/8	Down	0.0.0.0	0.0.0.0	Enabled	Inactive	00:11:22:33:44:55	1500
ge0/9	Down	0.0.0.0	0.0.0.0	Enabled	Inactive	00:11:22:33:44:55	1500
ge0/10	Down	0.0.0.0	0.0.0.0	Enabled	Inactive	00:11:22:33:44:55	1500

Figure 4.319 Routing > IP > Interface Summary

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. When viewing details about the routing settings for an interface, this field identifies the interface being viewed.

Item	Description
Status	Indicates whether the interface is capable of routing IP packets (Up) or cannot route packets (Down). For the status to be Up, the routing mode and administrative mode for the interface must be enabled. Additionally, the interface must have an IP address and be physically up (active link).
IP Address	The IP address of the interface.
Subnet Mask	The IP subnet mask for the interface (also known as the network mask or netmask). It defines the portion of the interface's IP address that is used to identify the attached network.
Admin Mode	The administrative mode of the interface, which is either Enabled or Disabled.
State	The state of the interface, which is either Active or Inactive. An interface is considered active if the link is up, and the interface is in a forwarding state.
MAC Address	The burned-in physical address of the interface. The format is six two-digit hexadecimal numbers separated by colons, for example 00:06:29:32:81:40.
IP MTU	The largest IP packet size the interface can transmit, in bytes. The IP Maximum Transmission Unit (MTU) is the maximum frame size minus the length of the Layer 2 header.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.
Details	Click Details to open a window and display additional information.
Add Loopback	Click Add Loopback to add the next available loopback interface. If maximum no of loopback interfaces are configured, the button is disabled.
Remove Loopback	Click Remove Loopback to remove the selected entries.

4.5.2.3 Interface Configuration




Use the Routing IP Interface Configuration page to configure the IP routing settings for each non-loopback interface.

To access this page, click **Routing > IP > Interface Configuration**.

Figure 4.320 Routing > IP > Interface Configuration

The following table describes the items in the previous figure.

Item	Description
Type	The type of interface that can be configured for routing: <ul style="list-style-type: none"> ■ Interface: Enables list of all non-loopback interfaces that can be configured for routing. ■ VLAN: Enables list of all VLANs that can be configured for routing.
VLAN	The menu contains all VLANs that can be configured for routing. To configure routing settings for a VLAN, select it from the menu and then configure the rest of the settings on the page.
Interface	The menu contains all non-loopback interfaces that can be configured for routing. To configure routing settings for an interface, select it from the menu and then configure the rest of the settings on the page.
Status	Indicates whether the interface is currently capable of routing IP packets (Up) or cannot route packets (Down). For the status to be Up, the routing mode and administrative mode for the interface must be enabled. Additionally, the interface must have an IP address and be physically up (active link).
Routing Mode	The administrative mode of IP routing on the interface.
Admin Mode	The administrative mode of the interface. If an interface is administratively disabled, it cannot forward traffic.
State	The state of the interface, which is either Active or Inactive. An interface is considered active if the link is up, and the interface is in a forwarding state.
Link Speed Data Rate	The physical link data rate of the interface.
IP Address Configuration Method	The method to use for configuring an IP address on the interface, which can be one of the following: <ul style="list-style-type: none"> ■ None: No address is to be configured. ■ Manual: The address is to be statically configured. When this option is selected you can specify the IP address and subnet mask in the available fields. ■ DHCP: The interface will attempt to acquire an IP address from a network DHCP server.
DHCP Client Identifier	The DHCP Client Identifier (Option 61) is used by DHCP clients to specify their unique identifier. DHCP servers use this value to index their database of address bindings. This value is expected to be unique for all clients in an administrative domain. The Client Identifier string will be displayed beside the check box once DHCP is enabled on the port on which the Client Identifier option is selected. This web page will need to be refreshed once this change is made.
IP Address	The IP address of the interface. This field can be configured only when the selected IP Address Configuration Method is Manual. If the method is DHCP, the interface attempts to lease an IP address from a DHCP server on the network, and the IP address appears in this field (read-only) after it is acquired. If this field is blank, the IP Address Configuration Method might be None, or the method might be DHCP and the interface is unable to lease an address.
Subnet Mask	The IP subnet mask for the interface (also known as the network mask or netmask). This field can be configured only when the selected IP Address Configuration Method is Manual.
MAC Address	The burned-in physical address of the interface. The format is six two-digit hexadecimal numbers separated by colons, for example 00:06:29:32:81:40.

Item	Description
IP MTU	The largest IP packet size the interface can transmit, in bytes. The IP Maximum Transmission Unit (MTU) is the maximum frame size minus the length of the Layer 2 header.
Bandwidth	The configured bandwidth on this interface. This setting communicates the speed of the interface to higher-level protocols.
Encapsulation Type	The link layer encapsulation type for packets transmitted from the interface, which can be either Ethernet or SNAP.
Forward Net Directed Broadcasts	Determines how the interface handles network-directed broadcast packets. A network-directed broadcast is a broadcast directed to a specific subnet. If this option is selected, network directed broadcasts are forwarded. If this option is clear, network directed broadcasts are dropped.
Proxy ARP	When this option is selected, proxy ARP is enabled, and the interface can respond to an ARP request for a host other than itself. An interface can act as an ARP proxy if it is aware of the destination and can route packets to the intended host, which is on a different subnet than the host that sent the ARP request.
Local Proxy ARP	When this option is selected, local proxy ARP is enabled, and the interface can respond to an ARP request for a host other than itself. Unlike proxy ARP, local proxy ARP allows the interface to respond to ARP requests for a host that is on the same subnet as the host that sent the ARP request. This feature is useful when a host is not permitted to reply to an ARP request from another host in the same subnet, for example when using the protected ports feature.
Destination Unreachables	When this option is selected, the interface is allowed to send ICMP Destination Unreachable message to a host if the intended destination cannot be reached for some reason. If this option is clear, the interface will not send ICMP Destination Unreachable messages to inform the host about the error in reaching the intended destination.
ICMP Redirects	When this option is selected, the interface is allowed to send ICMP Redirect messages. The device sends an ICMP Redirect message on an interface only if ICMP Redirects are enabled both globally and on the interface. An ICMP Redirect message notifies a host when a better route to a particular destination is available on the network segment.
Secondary IP Address	To add a secondary IP address on the interface, click  button in the header row and enter the address in the appropriate field in the Secondary IP Address Configuration window. You can add one or more secondary IP addresses to an interface only if the interface already has a primary IP address. To remove a configured secondary IP address, click  button associated with the entry to remove. To remove all configured secondary IP addresses, click  button in the header row.
Secondary Subnet Mask	The subnet mask associated with the secondary IP address. You configure this field in the Secondary IP Address Configuration window.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.5.2.4 Statistics

Use the Routing IP Loopback Configuration page to configure the IP routing settings for each loopback interface.

To access this page, click **Routing > IP > Statistics**.

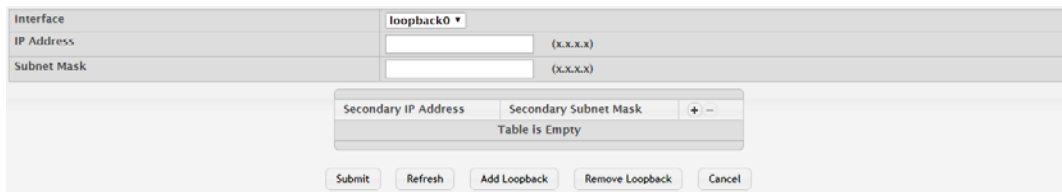


Figure 4.321 Routing > IP > Statistics

The following table describes the items in the previous figure.

Item	Description
Interface	The menu contains all loopback interfaces that can be configured for routing. To configure routing settings for an interface, select it from the menu and then configure the rest of the settings on the page.
IP Address	The IP address of the loopback interface.
Subnet Mask	The IP subnet mask for the interface (also known as the network mask or netmask).
Secondary IP Address	To add a secondary IP address on the interface, click the + (plus) symbol in the header row and enter the address in the appropriate field in the Secondary IP Address Configuration window. You can add one or more secondary IP addresses to an interface only if the interface already has a primary IP address. To remove a configured secondary IP address, click the: (minus) symbol associated with the entry to remove. To remove all configured secondary IP addresses, click the: (minus) symbol in the header row.
Secondary Subnet Mask	The subnet mask associated with the secondary IP address. This field is configurable in the Secondary IP Address Configuration window.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Add Loopback	Click Add Loopback to add the next available loopback interface. See the following procedure.
Remove Loopback	Click Remove Loopback to remove the selected entries.
Cancel	Click Cancel to restore default value.

4.5.2.5 Statistics

The Routing IP Statistics page displays information about the number and type of IP packets sent and received by all interfaces on the device. The statistics on this page are specified in RFC 1213.

To access this page, click **Routing > IP > Statistics**.

IpInReceives	31487
IpInHdrErrors	0
IpAddrErrors	0
IpFwdDatagrams	0
IpInUnknownProtos	0
IpInDiscards	0
IpInDelivers	31487
IpOutRequests	21177
IpOutDiscards	0
IpOutNoRoutes	6
IpReasmTimeout	0
IpReasmReqds	0
IpReasmOKs	0
IpReasmFails	0
IpFragOKs	0
IpFragFails	0
IpFragCreates	0
IpRoutingDiscards	0
IcmpInMsgs	3
IcmpInErrors	0
IcmpInDestUnreachs	3
IcmpInTimeExcds	0
IcmpInParmProbs	0
IcmpInSrcQuenchs	0
IcmpInRedirects	0
IcmpInEchos	0
IcmpInEchoReps	0
IcmpInTimeoetams	0

Figure 4.322 Routing > IP > Statistics

The following table describes the items in the previous figure.

Item	Description
IpInReceives	The total number of input datagrams received from all routing interfaces, including those datagrams received in error.
IpInHdrErrors	The number of input datagrams discarded due to errors in their IP headers, including bad checksums, version number mismatch, other format errors, time-to-live exceeded, errors discovered in processing their IP options, etc.
IpAddrErrors	The number of input datagrams discarded because the IP address in their IP header's destination field was not a valid address to be received at this entity. This count includes invalid addresses (e.g., 0.0.0.0) and addresses of unsupported classes (e.g., Class E). For entities which are not IP gateways and therefore do not forward datagrams, this counter includes datagrams discarded because the destination address was not a local address.
IpFwdDatagrams	The number of input datagrams for which this entity was not their final IP destination, as a result of which an attempt was made to find a route to forward them to that final destination. In entities which do not act as IP gateways, this counter will include only those packets which were Source-Routed via this entity, and the Source-Route option processing was successful.
IpInUnknownProtos	The number of locally-addressed datagrams received successfully but discarded because of an unknown or unsupported protocol.
IpInDiscards	The number of input IP datagrams for which no problems were encountered to prevent their continued processing, but which were discarded (e.g., for lack of buffer space). Note that this counter does not include any datagrams discarded while awaiting reassembly.
IpInDelivers	The total number of input datagrams successfully delivered to IP user-protocols (including ICMP).

Item	Description
IpOutRequests	The total number of IP datagrams which local IP user-protocols (including ICMP) supplied to IP in requests for transmission. Note that this counter does not include any datagrams counted in ipForwDatagrams.
IpOutDiscards	The number of output IP datagrams for which no problem was encountered to prevent their transmission to their destination, but which were discarded (e.g., for lack of buffer space). Note that this counter would include datagrams counted in ipForwDatagrams if any such packets met this (discretionary) discard criterion.
IpOutNoRoutes	The number of IP datagrams discarded because no route could be found to transmit them to their destination. Note that this counter includes any packets counted in ipForwDatagrams which meet this no-route criterion. Note that this includes any datagrams which a host cannot route because all of its default gateways are down.
IpReasmTimeout	The maximum number of seconds which received fragments are held while they are awaiting reassembly at this entity.
IpReasmReqds	The number of IP fragments received which needed to be reassembled at this entity.
IpReasmOKs	The number of IP datagrams successfully reassembled.
IpReasmFails	The number of failures detected by the IP reassembly algorithm (for whatever reason: timed out, errors, etc). Note that this is not necessarily a count of discarded IP fragments since some algorithms can lose track of the number of fragments by combining them as they are received.
IpFragOKs	The number of IP datagrams that have been successfully fragmented at this entity.
IpFragFails	The number of IP datagrams that have been discarded because they needed to be fragmented at this entity but could not be, e.g., because their Don't Fragment flag was set.
IpFragCreates	The number of IP datagram fragments that have been generated as a result of fragmentation at this entity.
IpRoutingDiscards	The number of routing entries which were chosen to be discarded even though they are valid. One possible reason for discarding such an entry could be to free-up buffer space for other routing entries.
IcmpInMsgs	The total number of ICMP messages which the entity received. Note that this counter includes all those counted by icmpInErrors.
IcmpInErrors	The number of ICMP messages which the entity received but determined as having ICMP-specific errors (bad ICMP checksums, bad length, etc.).
IcmpInDestUnreachs	The number of ICMP Destination Unreachable messages received.
IcmpInTimeExcds	The number of ICMP Time Exceeded messages received.
IcmpInParmProbs	The number of ICMP Parameter Problem messages received.
IcmpInSrcQuenchs	The number of ICMP Source Quench messages received.
IcmpInRedirects	The number of ICMP Redirect messages received.
IcmpInEchos	The number of ICMP Echo (request) messages received.
IcmpInEchoReps	The number of ICMP Echo Reply messages received.
IcmpInTimestamps	The number of ICMP Timestamp (request) messages received.
IcmpInTimestampReps	The number of ICMP Timestamp Reply messages received.
IcmpInAddrMasks	The number of ICMP Address Mask Request messages received.
IcmpInAddrMaskReps	The number of ICMP Address Mask Reply messages received.

Item	Description
IcmpOutMsgs	The total number of ICMP messages which this entity attempted to send. Note that this counter includes all those counted by icmpOutErrors.
IcmpOutErrors	The number of ICMP messages which this entity did not send due to problems discovered within ICMP, such as a lack of buffers. This value should not include errors discovered outside the ICMP layer such as the inability of IP to route the resultant datagram. In some implementations there may be no type of error that contributes to this counter's value.
IcmpOutDestUnreac hs	The number of ICMP Destination Unreachable messages sent.
IcmpOutTimeExcds	The number of ICMP Time Exceeded messages sent.
IcmpOutParmProbs	The number of ICMP Parameter Problem messages sent.
IcmpOutSrcQuenchs	The number of ICMP Source Quench messages sent.
IcmpOutRedirects	The number of ICMP Redirect messages sent. For a host, this object will always be zero, since hosts do not send redirects.
IcmpOutEchos	The number of ICMP Echo (request) messages sent.
IcmpOutEchoReps	The number of ICMP Echo Reply messages sent.
IcmpOutTimestamps	The number of ICMP Timestamp (request) messages.
IcmpOutTimestamp Reps	The number of ICMP Timestamp Reply messages sent.
IcmpOutAddrMasks	The number of ICMP Address Mask Request messages sent.
Refresh	Click Refresh to update the screen.

4.5.3 Router

4.5.3.1 Route Table

The Route Table Summary page collects routes from multiple sources: static routes and local routes. The route table manager may learn multiple routes to the same destination from multiple sources. The route table lists all routes. The best routes table displays only the most preferred route to each destination.

To access this page, click **Routing > Router > Route Table**.




Figure 4.323 Routing > Router > Route Table

The following table describes the items in the previous figure.

Item	Description
Network Address	The IP route prefix for the destination network.
Subnet Mask	The IP subnet mask (also known as the network mask or netmask) associated with the network address. It defines the portion of the IP address that is used to identify the attached network.

Item	Description
Protocol	Identifies which protocol created the route. A route can be created one of the following ways: <ul style="list-style-type: none"> ■ Dynamically learned through a supported routing protocol ■ Dynamically learned by being a directly-attached local route ■ Statically configured by an administrator ■ Configured as a default route by an administrator
Next Hop IP Address	The outgoing router IP address to use when forwarding traffic to the next router (if any) in the path towards the destination. The next router is always one of the adjacent neighbors or the IP address of the local interface for a directly-attached network.
Next Hop Interface	The outgoing interface to use when forwarding traffic to the destination. For a static reject route, the next hop is Null.
Best Route	Indicates whether the route is the preferred route to the network. If the field is blank, a better route to the same network exists in the routing table.
Refresh	Click Refresh to update the screen.

4.5.3.2 Configured Routes

Use the Configured Route Summary page to create and display static routes.

To access this page, click **Routing > Router > Configured Routes**.




Figure 4.324 Routing > Router > Configured Routes

The following table describes the items in the previous figure.

Item	Description
Network Address	The IP route prefix for the destination network. This IP address must contain only the network portion of the address and not the host bits. When adding a default route, this field is not available.
Subnet Mask	The IP subnet mask (also known as the network mask or netmask) associated with the network address. The subnet mask defines which portion of an IP address belongs to the network prefix, and which portion belongs to the host identifier. When adding a default route, this field is not available.
Next Hop IP Address	The outgoing router IP address to use when forwarding traffic to the next router (if any) in the path towards the destination. The next router is always one of the adjacent neighbors or the IP address of the local interface for a directly-attached network. When adding a static reject route, this field is not available because the packets are dropped rather than forwarded.
Next Hop Interface	The outgoing interface to use when forwarding traffic to the destination. For a static reject route, the next hop is Null.
Preference	The preference of the route. A lower preference value indicates a more preferred route. When the routing table has more than one route to the same network, the device selects the route with the best (lowest) route preference.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new route. See the following procedure.
Remove	Click Remove to remove the selected entries.

To add a new route:

Click **Routing > Router > Configured Routes > Add**.

Figure 4.325 Routing > Router > Configured Routes > Add

The following table describes the items in the previous figure.

Item	Description
Route Type	The type of route to configure, which is one of the following: <ul style="list-style-type: none"> ■ Default: The route the device uses to send a packet if the routing table does not contain a longer matching prefix for the packet's destination. The routing table can contain only one default route. ■ Static: A route that is manually added to the routing table by an administrator. ■ Static Reject: A route where packets that match the route are discarded instead of forwarded. The device might send an ICMP Destination Unreachable message.
Network Address	The IP route prefix for the destination network. This IP address must contain only the network portion of the address and not the host bits. When adding a default route, this field is not available.
Subnet Mask	The IP subnet mask (also known as the network mask or netmask) associated with the network address. The subnet mask defines which portion of an IP address belongs to the network prefix, and which portion belongs to the host identifier. When adding a default route, this field is not available.
Next Hop IP Address	The outgoing router IP address to use when forwarding traffic to the next router (if any) in the path towards the destination. The next router is always one of the adjacent neighbors or the IP address of the local interface for a directly-attached network. When adding a static reject route, this field is not available because the packets are dropped rather than forwarded.
Preference	The preference of the route. A lower preference value indicates a more preferred route. When the routing table has more than one route to the same network, the device selects the route with the best (lowest) route preference.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.5.3.3 Summary

Use the IP Route Summary page to view the statistics of static routes.

To access this page, click **Routing > Router > Summary**.

Route Types	
Connected Routes	0
Static Routes	0
RIP Routes	0
OSPF Routes	0
Intra Area Routes	0
Inter Area Routes	0
External Type-1 Routes	0
External Type-2 Routes	0
Reject Routes	0
Total Routes	0
Route Table Counters	
Best Routes (High)	0 (0)
Alternate Routes	0
Route Adds	0
Route Modifies	0
Route Deletes	0
Unresolved Route Adds	0
Invalid Route Adds	0
Failed Route Adds	0
Reserved Locals	0
Unique Next Hops (High)	0 (0)
Next Hop Groups (High)	0 (0)
ECMP Groups (High)	0 (0)
ECMP Routes	0
Truncated ECMP Routes	0
ECMP Retries	0

Refresh Clear Counters

Figure 4.326 Routing > Router > Summary

The following table describes the items in the previous figure.

Item	Description
Route Types	
Connected Routes	The total number of connected routes in the IP routing table.
Static Routes	The total number of static routes in the IP routing table.
RIP Routes	The total number of routes installed by the RIP protocol.
OSPF Routes	The total number of routes installed by the OSPF protocol. <ul style="list-style-type: none"> ■ Intra Area Routes: The total number of intra-area routes installed by the OSPF protocol. ■ Inter Area Routes: The total number of inter-area routes installed by the OSPF protocol. ■ External Type-1 Routes: The total number of external type-1 routes installed by the OSPF protocol. ■ External Type-2 Routes: The total number of external type-2 routes installed by the OSPF protocol.
Reject Routes	The total number of reject routes installed by all protocols.
Total Routes	The total number of routes in the routing table.
Route Table Counters	
Best Routes (High)	The number of best routes currently in the routing table. This number only counts the best route to each destination.
Alternate Routes	The number of alternate routes currently in the routing table. An alternate route is a route that was not selected as the best route to its destination.
Route Adds	The number of routes that have been added to the routing table.
Route Modifies	The number of routes that have been changed after they were initially added to the routing table.
Route Deletes	The number of routes that have been deleted from the routing table.

Item	Description
Unresolved Route Adds	The number of route adds that failed because none of the route's next hops were on a local subnet. Note that static routes can fail to be added to the routing table at startup because the routing interfaces are not yet up. This counter gets incremented in this case. The static routes are added to the routing table when the routing interfaces come up.
Invalid Route Adds	The number of routes that failed to be added to the routing table because the route was invalid. A log message is written for each of these failures.
Failed Route Adds	The number of routes that failed to be added to the routing table because of a resource limitation in the routing table.
Reserved Locals	The number of routing table entries reserved for a local subnet on a routing interface that is down. Space for local routes is always reserved so that local routes can be installed when a routing interface bounces.
Unique Next Hops (High)	The number of distinct next hops used among all routes currently in the routing table. These include local interfaces for local routes and neighbors for indirect routes.
Next Hop Groups (High)	The current number of next hop groups in use by one or more routes. Each next hop group includes one or more next hops.
ECMP Groups (High)	The number of next hop groups with multiple next hops.
ECMP Routes	The number of routes with multiple next hops currently in the routing table.
Truncated ECMP Routes	The number of ECMP routes that are currently installed in the forwarding table with just one next hop. The forwarding table may limit the number of ECMP routes or the number of ECMP groups. When an ECMP route cannot be installed because such a limit is reached, the route is installed with a single next hop.
ECMP Retries	The number of ECMP routes that have been installed in the forwarding table after initially being installed with a single next hop.
Refresh	Click Refresh to update the screen.
Clear Counters	Click Clear Counters to reset all counters to zero.

4.5.3.4 ECMP Group

The ECMP Groups Summary page displays all current ECMP groups in the IPv4 routing table. An ECMP group is a set of two or more next hops used in one or more routes.

To access this page, click **Routing > Router > ECMP Group**.




Figure 4.327 Routing > Router > ECMP Group

The following table describes the items in the previous figure.

Item	Description
ECMP Group	The ECMP group number associated with the rest of the data in the row. The device assigns an arbitrary ECMP group number from 1 to n to identify the ECMP group.
Number Of Next Hops	The number of next hops in the group.

Item	Description
Route(s) Count	The number of routes that use the set of next hops.
Next Hops	The IPv4 address and outgoing interface of each next hop in the group.
Refresh	Click Refresh to update the screen.

4.5.4 IPv6

4.5.4.1 Configuration

Use the IPv6 Global Configuration page to configure global IPv6 routing settings on the device. IPv6 routing provides a means of transmitting IPv6 packets between subnets on the network. IPv6 routing configuration is necessary only if the device is used as a Layer 3 device that routes IPv6 packets between subnets. IPv6 is the next generation of the Internet Protocol. With 128-bit addresses, versus 32-bit addresses for IPv4, IPv6 solves the address depletion issues seen with IPv4 and removes the requirement for Network Address Translation (NAT), which is used in IPv4 networks to reduce the number of globally unique IP addresses required for a given network.

To access this page, click **Routing > IPv6 > Configuration**.

The screenshot shows the IPv6 Configuration page with the following settings:

Item	Value	Range
IPv6 Unicast Routing Mode	Disable (selected)	Enable
IPv6 Neighbors Dynamic Renew	<input type="checkbox"/>	
IPv6 Hop Limit	64	(1 to 255)
IPv6 Unresolved Packets Rate Limit (pps)	1024	(50 to 1024)
NUD Maximum Unicast Solicits	3	(3 to 10)
NUD Maximum Multicast Solicits	3	(3 to 255)
NUD Back-off Multiple	1	(1 to 5)
ICMPv6 Rate Limit Error Interval (Msecs)	1000	(0 to 2147483647)
ICMPv6 Rate Limit Burst Size	100	(1 to 200)
Static Route Preference	1	(1 to 255)
Local Route Preference	0	

Buttons: Submit, Refresh, Cancel

Figure 4.328 Routing > IPv6 > Configuration

The following table describes the items in the previous figure.

Item	Description
IPv6 Unicast Routing Mode	The administrative mode of IPv6 routing on the device. The options are as follows: <ul style="list-style-type: none"> ■ Enable: The device can act as a Layer 3 device by routing IPv6 packets between interfaces configured for IPv6 routing. ■ Disable: The device does not support IPv6 routing.
IPv6 Neighbors Dynamic Renew	Select this option to enable dynamic renewal mode for the periodic Neighbor Unreachability Detection (NUD) run on the existing IPv6 neighbor entries in the IPv6 neighbor cache. If NUD attempts to communicate with IPv6 neighbors and no response is received after the maximum number of solicits is reached, its entry is removed from the cache.
IPv6 Hop Limit	The unicast hop count used in IPv6 packets originated by the device. This value is also included in router advertisements.
IPv6 Unresolved Packets Rate Limit	The rate in packets-per-second for the number of IPv6 data packets trapped to the CPU when the packet fails to be forwarded in the hardware due to the unresolved hardware address of the destined IPv6 node.
NUD Maximum Unicast Solicits	The maximum number of unicast neighbor solicitations sent during NUD before switching to multicast neighbor solicitations.
NUD Maximum Multicast Solicits	The maximum number of multicast neighbor solicitations sent during NUD when a neighbor is in the UNREACHABLE state.

Item	Description
NUD Back-off Multiple	The exponential backoff multiplier to be used in the calculation of the next timeout value for neighbor solicitation transmission during NUD following the exponential backoff algorithm.
ICMPv6 Rate Limit Error Interval	The maximum burst interval for ICMPv6 error messages transmitted by the device. The rate limit for ICMPv6 error messages is configured as a token bucket. The ICMPv6 Rate Limit Error Interval specifies how often the token bucket is initialized with tokens of the size configured in the ICMPv6 Rate Limit Burst Size field.
ICMPv6 Rate Limit Burst Size	The number of ICMPv6 error messages that can be sent during the burst interval configured in the ICMPv6 Rate Limit Error Interval field.
Static Route Preference	The default distance (preference) for static IPv6 routes. Lower route-distance values are preferred when determining the best route. The value configured for Static Route Preference is used when using the CLI to configure a static route and no preference is specified. Changing the Static Route Preference does not update the preference of existing static routes.
Local Route Preference	The default distance (preference) for local IPv6 routes.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.5.4.2 Interface Summary

The IPv6 Interface Summary page shows summary information about the IPv6 routing configuration for all interfaces.

To access this page, click **Routing > IPv6 > Interface Summary**.

The screenshot shows a web interface for IPv6 Interface Summary. At the top, it says 'Display 10 rows' and 'Showing 1 to 10 of 29 entries'. There is a search filter box. Below is a table with columns: Interface, Operational Status, IPv6 Mode, Routing Mode, Admin Mode, IPv6 Prefix, Prefix Length, and State. The table lists interfaces 0/1 through 0/10. All Operational Status and Admin Mode are 'Enabled', while IPv6 Mode and Routing Mode are 'Disabled'. Below the table are navigation buttons: First, Previous, 1, 2, 3, Next, Last. At the bottom are buttons for Refresh, Edit, Details, and Add Loopback.

Interface	Operational Status	IPv6 Mode	Routing Mode	Admin Mode	IPv6 Prefix	Prefix Length	State
0/1	Disabled	Disabled	Disabled	Enabled			
0/2	Disabled	Disabled	Disabled	Enabled			
0/3	Disabled	Disabled	Disabled	Enabled			
0/4	Disabled	Disabled	Disabled	Enabled			
0/5	Disabled	Disabled	Disabled	Enabled			
0/6	Disabled	Disabled	Disabled	Enabled			
0/7	Disabled	Disabled	Disabled	Enabled			
0/8	Disabled	Disabled	Disabled	Enabled			
0/9	Disabled	Disabled	Disabled	Enabled			
0/10	Disabled	Disabled	Disabled	Enabled			

Figure 4.329 Routing > IPv6 > Interface Summary

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. When viewing details about the routing settings for an interface, this field identifies the interface being viewed.
Operational Status	The IPv6 routing operational mode on the interface. The operational mode is enabled under the conditions in the following list; otherwise, the mode is disabled: <ul style="list-style-type: none"> ■ The IPv6 mode is enabled on the interface. ■ The routing mode is enabled on the interface. ■ The administrative mode is enabled on the interface. ■ The link is up.

Item	Description
IPv6 Mode	The IPv6 mode on the interface. When IPv6 mode is enabled, the interface is capable of IPv6 operation without a global address. In this case, an EUI-64 based link-local address is used.
Routing Mode	Indicates whether Layer 3 routing is administratively enabled or disabled on the interface.
Admin Mode	The administrative mode on the interface.
IPv6 Prefix	The IPv6 routing prefix dynamically or manually configured on the interface.
Prefix Length	The number of bits used for the IPv6 prefix.
State	The state of the IPV6 address. The state is TENT if routing is disabled or Duplicate Address Detection (DAD) fails. The state is Active if the interface is active and DAD is successful. Otherwise, the state is Inactive.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.
Details	Click Details to open the IPv6 Interface Details window.
Add Loopback	Click Add Loopback to add the next available loopback interface.

4.5.4.3 Interface Configuration

Use the IPv6 Interface Configuration page to configure the IPv6 routing settings for each non-loopback interface.

To access this page, click **Routing > IPv6 > Interface Configuration**.

Figure 4.330 Routing > IPv6 > Interface Configuration

The following table describes the items in the previous figure.

Item	Description
Type	The type of interface that can be configured for IPv6 routing: <ul style="list-style-type: none"> Interface: Enables list of all non-loopback interfaces that can be configured for IPv6 routing. VLAN: Enables list of all VLANs that can be configured for IPv6 routing.
VLAN	The menu contains all VLANs that can be configured for IPv6 routing. To configure routing settings for a VLAN, select it from the menu and then configure the rest of the settings on the page.

Item	Description
Interface	The menu contains all non-loopback interfaces that can be configured for IPv6 routing. To configure routing settings for an interface, select it from the menu and then configure the rest of the settings on the page.
Operational Status	The IPv6 routing operational mode on the interface. The operational mode is enabled under the conditions in the following list; otherwise, the mode is disabled: <ul style="list-style-type: none"> ■ The IPv6 mode is enabled on the interface. ■ The routing mode is enabled on the interface. ■ The administrative mode is enabled on the interface. ■ The link is up.
Link Local Prefix	The autoconfigured link-local address which is: <ul style="list-style-type: none"> ■ Allocated from part of the IPv6 unicast address space ■ Not visible off the local link ■ Not globally unique
Link Local Prefix Length	The number of bits used for the prefix of the link-local IPv6 address.
Link Local Status	The status of the IPV6 link local address. The status is TENT if routing is disabled or Duplicate Address Detection (DAD) fails. The state is Active if the interface is active and DAD is successful. Otherwise, the state is Inactive.
Routing Mode	The administrative mode for Layer 3 routing on the interface.
IPv6 Mode	The administrative mode for IPv6 on the interface. When IPv6 mode is enabled, the interface is capable of IPv6 operation without a global address. In this case, an EUI-64 based link-local address is used.
Admin Mode	The administrative mode of the interface. If an interface is administratively disabled, it will not forward traffic.
DHCPv6 Client Mode	The administrative mode of the DHCPv6 client on the interface. An interface can acquire an IPv6 address by communicating with a network DHCPv6 server (stateful configuration). An interface can also obtain an IPv6 address through stateless address autoconfiguration or static configuration.
Stateless Address AutoConfig	When this option is selected, the interface can generate its own IPv6 address by using local interface information and prefix information advertised by routers.
Interface Maximum Transmit Unit	The largest IPv6 packet size the interface can transmit, in bytes. To change the MTU value, click the Edit icon to the right of the field. To reset the MTU to the default value, click the Reset icon.
Router Duplicate Address Detection Transmits	The number of duplicate address detection probes the interface transmits while doing neighbor discovery.
Router Advertisement NS Interval	The interval between router advertisements for advertised neighbor solicitations. To change the interval, click the Edit icon to the right of the field. To reset the interval to the default value, click the Reset icon.
Router Lifetime Interval	The value that is placed in the Router Lifetime field of the router advertisements sent from the interface.
Router Advertisement Reachable Time	The value that is placed in the Reachable Time field of the router advertisements. The amount of time to consider a neighbor reachable after neighbor discovery confirmation.
Router Advertisement Interval	The transmission interval between router advertisements messages sent by the interface.

Item	Description
Router Advertisement Managed Config	When this option is selected, the Managed Address Configuration flag is set in router advertisements, indicating that the interface should use stateful autoconfiguration (DHCPv6) to obtain addresses.
Router Advertisement Other Config	When this option is selected, the Other Stateful Configuration flag is set in router advertisements, indicating that the interface should use stateful autoconfiguration for information other than addresses.
Router Advertisement Suppress	When this option is selected, the interface does not transmit router advertisements.
IPv6 Destination Unreachable Messages	When this option is selected, the interface can send ICMPv6 Destination Unreachable messages back to the source device if a datagram is undeliverable.
ICMPv6 Redirects	When this option is selected, the interface is allowed to send ICMPv6 Redirect messages. An ICMPv6 Redirect message notifies a host when a better route to a particular destination is available on the network segment.
IPv6 Hop Limit Unspecified	When this option is selected, the device can send Router Advertisements on this interface with an unspecified (0) current hop limit value. This will tell the hosts on the link to ignore the hop limit from this device.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.5.4.4 Loopback Configuration

Use the IPv6 Loopback Configuration page to configure the IPv6 routing settings for each loopback interface. A loopback interface is a logical interface that is always up (as long as it is administratively enabled) and, because it cannot go down, allows the device to have a stable IPv6 address that other network nodes and protocols can use to reach the device. The loopback can provide the source address for sent packets. The loopback interface does not behave like a network switching port. Specifically, there are no neighbors on a loopback interface; it is a pseudodevice for assigning local addresses so that the other Layer 3 hosts can communicate with the device by using the loopback IPv6 address.

To access this page, click **Routing > IPv6 > Loopback Configuration**.

Figure 4.331 Routing > IPv6 > Loopback Configuration

The following table describes the items in the previous figure.

Item	Description
Interface	The menu contains all loopback interfaces that can be configured for IPv6 routing. To configure routing settings for an interface, select it from the menu and then configure the rest of the settings on the page.
Operational Status	The operational status of the loopback interface. To be operational, both the IPv6 mode and administrative mode must be enabled.

Item	Description
Link Local Prefix	The autoconfigured link-local address which is: <ul style="list-style-type: none"> Allocated from part of the IPv6 unicast address space Not visible off the local link Not globally unique
Link Local Prefix Length	The number of bits used for the prefix of the link-local IPv6 address.
IPv6 Mode	The IPv6 mode on the loopback interface. When IPv6 mode is enabled, the interface is capable of IPv6 operation without a global address. In this case, an EUI-64 based link-local address is used.
Admin Mode	The administrative mode of the loopback interface.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.5.4.5 Global Addresses

The IPv6 Global Address Table page shows information about all global IPv6 addresses configured on all interfaces on the device. From this page, you can also remove a configured IPv6 address from an interface.

To access this page, click **Routing > IPv6 > Global Addresses**.




Figure 4.332 Routing > IPv6 > Global Addresses

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row.
IPv6 Prefix	The IPv6 routing prefix dynamically or manually configured on the interface. This page does not show information about link-local addresses.
Prefix Length	The number of bits used for the IPv6 prefix.
State	The link state, which is either Active or Inactive.
Valid Lifetime	The value, in seconds, to be placed in the Valid Lifetime field of the Prefix Information option in a router advertisement. The prefix is valid for on-link determination for this length of time. Hosts that generate an address from this prefix using stateless address auto-configuration can use those addresses for this length of time. An auto-configured address older than the preferred lifetime but younger than the valid lifetime are considered deprecated addresses. As defined by RFC 2462, a deprecated address is "An address assigned to an interface whose use is discouraged, but not forbidden. A deprecated address should no longer be used as a source address in new communications, but packets sent from or to deprecated addresses are delivered as expected. A deprecated address may continue to be used as a source address in communications where switching to a preferred address causes hardship to a specific upper-layer activity (e.g., an existing TCP connection)." If you specify the maximum value, an autoconfigured address may remain preferred indefinitely.

Item	Description
Preferred Lifetime	The value, in seconds, to be placed in the Preferred Lifetime in the Prefix Information option in a router advertisement. Addresses generated from a prefix using stateless address autoconfiguration remain preferred for this length of time. If you specify the maximum value, an autoconfigured address may remain preferred indefinitely.
Onlink Flag	The state of the on-link flag in the IPv6 prefix. When enabled, the prefix can be used for on-link determination by other hosts with IPv6 addresses within this prefix.
Autonomous Flag	The state of the autonomous flag in the IPv6 prefix. When enabled, the prefix can be used for autonomous address configuration by other hosts (in combination with an interface identifier on the other hosts).
Submit	Click Submit to save the values and update the screen.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

4.5.4.6 Address Configuration

Use the IPv6 Global Address Configuration page to statically configure a global IPv6 addresses on an interface. A global address is globally routable and is recognized outside of the local network. To configure an IPv6 address on an interface that already has an IPv6 address, click Add.

To access this page, click **Routing > IPv6 > Address Configuration**.

Figure 4.333 Routing > IPv6 > Address Configuration

The following table describes the items in the previous figure.

Item	Description
Interface	The menu contains all interfaces valid for routing that exist on the system. Select an interface to configure an address.
IPv6 Prefix	The global IPv6 routing prefix dynamically or manually configured on the interface. If the interface has more than one IPv6 address, use the menu to select the address to view or update.
Prefix Length	The number of bits used for the IPv6 prefix.
Valid Lifetime	The value, in seconds, to be placed in the Valid Lifetime field of the Prefix Information option in a router advertisement. The prefix is valid for on-link determination for this length of time. Hosts that generate an address from this prefix using stateless address auto-configuration can use those addresses for this length of time. An auto-configured address older than the preferred lifetime but younger than the valid lifetime are considered deprecated addresses. If you specify the maximum value, an autoconfigured address may remain preferred indefinitely.
Preferred Lifetime	The value, in seconds, to be placed in the Preferred Lifetime in the Prefix Information option in a router advertisement. Addresses generated from a prefix using stateless address autoconfiguration remain preferred for this length of time. If you specify the maximum value, an autoconfigured address may remain preferred indefinitely.

Item	Description
Onlink Flag	The state of the on-link flag in the IPv6 prefix. When enabled, the prefix that can be used for on-link determination by other hosts with IPv6 addresses within this prefix.
Autonomous Flag	The state of the autonomous flag in the IPv6 prefix. When enabled, the prefix can be used for autonomous address configuration by other hosts (in combination with an interface identifier on the other hosts).
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.
Add	Click Add to configure an IPv6 address on an interface that already has an IPv6 address.

4.5.4.7 Statistics

The IPv6 Statistics page displays summary statistics about the IPv6 datagrams each interface sends and receives, successfully or unsuccessfully. It also displays summary statistics about the ICMPv6 messages each interface sends and receives. To view more information about the types of datagrams and IPv6 messages an interface has sent and received, select the interface with the information to view and click Details. You are redirected to the IPv6 Detailed Statistics page for the selected interface.

To access this page, click **Routing > IPv6 > Statistics**.

Interface	0/1
Total Datagrams Received	0
Datagrams Forwarded	0
Total ICMPv6 Messages Received	0
ICMPv6 Messages With Errors Received	0
Total ICMPv6 Messages Transmitted	0
ICMPv6 Duplicate Address Detects	0

Figure 4.334 Routing > IPv6 > Statistics

The following table describes the items in the previous figure.

Item	Description
Interface	The menu contains all interfaces valid for routing that exist on the system. Select an interface to view its IPv6 statistics.
Total Datagrams Received	The total number of input datagrams received by the interface, including those received in error.
Datagrams Forwarded	The number of output datagrams that this entity received and forwarded toward their final destinations. In entities that do not act as IPv6 routers, this counter will include only those packets which were source-routed via this entity, and the source-route processing was successful. Note that for a successfully forwarded datagram, the counter of the outgoing interface is incremented.
Total ICMPv6 Messages Received	The total number of ICMPv6 messages received by the interface, which includes all those counted by ipv6IcmpInErrors. Note that this interface is the interface to which the ICMPv6 messages were addressed, which may not be necessarily the input interface for the messages.
ICMPv6 Messages With Errors Received	The number of ICMPv6 messages that the interface received but were determined to have ICMPv6-specific errors (bad ICMPv6 checksums, bad length, etc.)
Total ICMPv6 Messages Transmitted	The total number of ICMPv6 messages that this interface attempted to send. Note that this counter includes all those counted by icmpOutErrors.

Item	Description
ICMPv6 Duplicate Address Detects	The number of duplicate IPv6 addresses detected by the interface.
Refresh	Click Refresh to update the screen.
Details	Click Details to open the Group Entry Details window.

4.5.4.8 Detailed Statistics

To access this page, click **Routing > IPv6 > Detailed Statistics**.

Interface	
0/1	
IPv6 Statistics	
Total Datagrams Received	0
Received Datagrams Locally Delivered	0
Received Datagrams Discarded Due To Header Errors	0
Received Datagrams Discarded Due To MTU	0
Received Datagrams Discarded Due To No Route	0
Received Datagrams With Unknown Protocol	0
Received Datagrams Discarded Due To Invalid Address	0
Received Datagrams Discarded Due To Truncated Data	0
Received Datagrams Discarded Other	0
Received Datagrams Reassembly Required	0
Datagrams Successfully Reassembled	0
Datagrams Failed To Reassemble	0
Datagrams Forwarded	0
Datagrams Locally Transmitted	0
Datagrams Transmit Failed	0
Datagrams Fragments Created	0
Datagrams Failed To Fragment	0
Datagrams Successfully Fragmented	0
Multicast Datagrams Received	0
Multicast Datagrams Transmitted	0
ICMPv6 Statistics	
Total ICMPv6 Messages Received	0
ICMPv6 Messages With Errors Received	0
ICMPv6 Destination Unreachable Messages Received	0

Figure 4.335 Routing > IPv6 > Detailed Statistics

The following table describes the items in the previous figure.

Item	Description
Interface	The menu contains all interfaces valid for routing that exist on the system. Select an interface to view its IPv6 statistics.
IPv6 Statistics	
Total Datagrams Received	The total number of input datagrams received by the interface, including those received in error.
Received Datagrams Locally Delivered	The total number of datagrams successfully delivered to IPv6 user-protocols (including ICMP). This counter is incremented at the interface to which these datagrams were addressed which might not be necessarily the input interface for some of the datagrams.
Received Datagrams Discarded Due To Header Errors	The number of input datagrams discarded due to errors in their IPv6 headers, including version number mismatch, other format errors, hop count exceeded, errors discovered in processing their IPv6 options, etc.
Received Datagrams Discarded Due To MTU	The number of input datagrams that could not be forwarded because their size exceeded the link MTU of outgoing interface.
Received Datagrams Discarded Due To No Route	The number of input datagrams discarded because no route could be found to transmit them to their destination.
Received Datagrams With Unknown Protocol	The number of locally-addressed datagrams received successfully but discarded because of an unknown or unsupported protocol. This counter is incremented at the interface to which these datagrams were addressed which might not be necessarily the input interface for some of the datagrams.

Item	Description
Received Datagrams Discarded Due To Invalid Address	The number of input datagrams discarded because the IPv6 address in their IPv6 header's destination field was not a valid address to be received at this entity. This count includes invalid addresses, e.g., ::0, and unsupported addresses, e.g., addresses with unallocated prefixes. For entities which are not IPv6 routers and therefore do not forward datagrams, this counter includes datagrams discarded because the destination address was not a local address.
Received Datagrams Discarded Due To Truncated Data	The number of input datagrams discarded because datagram frame didn't carry enough data.
Received Datagrams Discarded Other	The number of input IPv6 datagrams for which no problems were encountered to prevent their continued processing, but which were discarded (e.g., for lack of buffer space). Note that this counter does not include any datagrams discarded while awaiting re-assembly.
Received Datagrams Reassembly Required	The number of IPv6 fragments received which needed to be reassembled at this interface. Note that this counter is incremented at the interface to which these fragments were addressed which might not be necessarily the input interface for some of the fragments.
Datagrams Successfully Reassembled	The number of IPv6 datagrams successfully reassembled. Note that this counter is incremented at the interface to which these datagrams were addressed which might not be necessarily the input interface for some of the fragments.
Datagrams Failed To Reassemble	The number of failures detected by the IPv6 reassembly algorithm (for whatever reason: timed out, errors, etc.). Note that this is not necessarily a count of discarded IPv6 fragments since some algorithms (notably the algorithm in RFC 815) can lose track of the number of fragments by combining them as they are received. This counter is incremented at the interface to which these fragments were addressed which might not be necessarily the input interface for some of the fragments.
Datagrams Forwarded	The number of output datagrams that this entity received and forwarded toward their final destinations. In entities that do not act as IPv6 routers, this counter will include only those packets which were source-routed via this entity, and the source-route processing was successful. Note that for a successfully forwarded datagram, the counter of the outgoing interface is incremented.
Datagrams Locally Transmitted	The number of datagrams which this entity has successfully transmitted from this output interface.
Datagrams Transmit Failed	The number of datagrams which this entity failed to transmit successfully.
Datagrams Fragments Created	The number of IPv6 datagrams that have been successfully fragmented at this output interface.
Datagrams Failed To Fragment	The number of output datagrams that could not be fragmented at this interface.
Datagrams Successfully Fragmented	The number of output datagram fragments that have been generated as a result of fragmentation at this output interface.
Multicast Datagrams Received	The number of multicast packets received by the interface.
Multicast Datagrams Transmitted	The number of multicast packets transmitted by the interface.
ICMPv6 Statistics	

Item	Description
Total ICMPv6 Messages Received	The total number of ICMPv6 messages received by the interface, which includes all those counted by ipv6IcmpInErrors. Note that this interface is the interface to which the ICMPv6 messages were addressed, which may not be necessarily the input interface for the messages.
ICMPv6 Messages With Errors Received	The number of ICMPv6 messages that the interface received but were determined to have ICMPv6-specific errors (bad ICMPv6 checksums, bad length, etc.)
ICMPv6 Destination Unreachable Messages Received	The number of ICMPv6 Destination Unreachable messages received by the interface.
ICMPv6 Messages Prohibited Administratively Received	The number of ICMPv6 destination unreachable/communication administratively prohibited messages received by the interface.
ICMPv6 Time Exceeded Messages Received	The number of ICMPv6 Time Exceeded messages received by the interface.
ICMPv6 Parameter Problem Messages Received	The number of ICMPv6 Parameter Problem messages received by the interface.
ICMPv6 Packet Too Big Messages Received	The number of ICMPv6 Packet Too Big messages received by the interface.
ICMPv6 Echo Request Messages Received	The number of ICMPv6 Echo (request) messages received by the interface.
ICMPv6 Echo Reply Messages Received'	The number of ICMPv6 Echo Reply messages received by the interface.
ICMPv6 Router Solicit Messages Received	The number of ICMPv6 Neighbor Solicitation messages received by the interface.
ICMPv6 Router Advertisement Messages Received	The number of ICMPv6 Router Advertisement messages received by the interface.
ICMPv6 Neighbor Solicit Messages Received	The number of ICMPv6 Neighbor Solicitation messages received by the interface.
ICMPv6 Neighbor Advertisement Messages Received	The number of ICMPv6 Neighbor Advertisement messages received by the interface.
ICMPv6 Redirect Messages Received	The number of Redirect messages received.
ICMPv6 Group Membership Query Messages Received	The number of ICMPv6 Group Membership Query messages received.
ICMPv6 Group Membership Response Messages Received	The number of ICMPv6 Group Membership Response messages received.
ICMPv6 Group Membership Response Messages Received	The number of ICMPv6 Group Membership Reduction messages received.

Item	Description
Total ICMPv6 Messages Transmitted	The total number of ICMPv6 messages which this interface attempted to send. Note that this counter includes all those counted by icmpOutErrors.
ICMPv6 Messages Not Transmitted Due To Error	The number of ICMPv6 messages which this interface did not send due to problems discovered within ICMPv6 such as a lack of buffers. This value should not include errors discovered outside the ICMPv6 layer such as the inability of IPv6 to route the resultant datagram. In some implementations there may be no types of error which contribute to this counter's value.
ICMPv6 Destination Unreachable Messages Transmitted	The number of ICMPv6 Destination Unreachable Messages sent by the interface.
ICMPv6 Messages Prohibited Administratively Transmitted	The number of ICMPv6 destination unreachable/communication administratively prohibited messages sent.
ICMPv6 Time Exceeded Messages Transmitted	The number of ICMPv6 Time Exceeded messages sent by the interface.
ICMPv6 Parameter Problem Messages Transmitted	The number of ICMPv6 Parameter Problem messages sent by the interface.
ICMPv6 Packet Too Big Messages Transmitted	The number of ICMPv6 Packet Too Big messages sent by the interface.
ICMPv6 Echo Request Messages Transmitted	The number of ICMPv6 Echo (request) messages sent by the interface.
ICMPv6 Router Solicit Messages Transmitted	The number of ICMPv6 Router Solicitation messages sent by the interface.
ICMPv6 Router Advertisement Messages Transmitted	The number of ICMPv6 Router Advertisement messages sent by the interface.
ICMPv6 Neighbor Solicit Messages Transmitted	The number of ICMPv6 Neighbor Solicitation messages sent by the interface.
ICMPv6 Neighbor Advertisement Messages Transmitted	The number of ICMPv6 Neighbor Advertisement messages sent by the interface.
ICMPv6 Redirect Messages Transmitted	The number of Redirect messages sent.
ICMPv6 Group Membership Query Messages Transmitted	The number of ICMPv6 Group Membership Query messages sent.
ICMPv6 Group Membership Response Messages Transmitted	The number of ICMPv6 Group Membership Response messages sent.

Item	Description
ICMPv6 Group Membership Reduction Messages Transmitted	The number of ICMPv6 Group Membership Reduction messages sent.
ICMPv6 Duplicate Address Detects	The number of duplicate IPv6 addresses detected by the interface.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.5.4.9 Neighbor Table

The IPv6 Neighbor Table page displays the IPv6 neighbor entries in the local IPv6 neighbor cache. Neighbors are discovered by using the Neighbor Discovery Protocol via ICMPv6 messages on active IPv6 interfaces.

To access this page, click **Routing > IPv6 > Neighbor Table**.




Figure 4.336 Routing > IPv6 > Neighbor Table

The following table describes the items in the previous figure.

Item	Description
Interface	The local interface on which the neighbor was discovered.
IPv6 Address	The IPv6 prefix and prefix length of the neighbor interface.
MAC Address	The MAC address associated with the neighbor interface. If the MAC address is all zeros, the entry is a Negative NDP entry. A Negative NDP entry is added to the table when the device sends a Neighbor Solicitation Request, but it has not yet been resolved. If the request is resolved and the neighbor is reachable, its valid MAC address replaces the null address. If the request times out, the entry is removed.
Type	The type of the neighbor entry, which is one of the following: <ul style="list-style-type: none"> ■ Static: The neighbor entry is manually configured. ■ Dynamic: The neighbor entry is dynamically resolved. ■ Local: The neighbor entry is a local entry. ■ Other: The neighbor entry is an unknown entry.
Is Router	Indicates whether the neighbor is a router. If the neighbor is a router, the value is TRUE. If the neighbor is not a router, the value is FALSE.

Item	Description
Neighbor State	<p>Specifies the state of the neighbor cache entry. Dynamic entries in the IPv6 neighbor discovery cache can be one of the following:</p> <ul style="list-style-type: none"> ■ Incmp: Address resolution is being performed on the entry. A neighbor solicitation message has been sent to the solicited-node multicast address of the target, but the corresponding neighbor advertisement message has not yet been received. ■ Reach: Positive confirmation was received within the last Reachable Time milliseconds that the forward path to the neighbor was functioning properly. While in REACH state, the device takes no special action as packets are sent. ■ Stale: More than ReachableTime milliseconds have elapsed since the last positive confirmation was received that the forward path was functioning properly. While in STALE state, the device takes no action until a packet is sent. ■ Delay: More than ReachableTime milliseconds have elapsed since the last positive confirmation was received that the forward path was functioning properly. A packet was sent within the last DELAY_FIRST_PROBE_TIME seconds. If no reachability confirmation is received within DELAY_FIRST_PROBE_TIME seconds of entering the DELAY state, send a neighbor solicitation message and change the state to PROBE. ■ Probe: A reachability confirmation is actively sought by resending neighbor solicitation messages every RetransTimer milliseconds until a reachability confirmation is received.
Last Updated	The amount of time that has passed since the address was confirmed to be reachable.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new IPv6 Neighbor. See the following procedure.
Remove	Click Remove to remove the selected entries.

To add a new IPv6 Neighbor:

Click **Routing > IPv6 > Neighbor Table > Add**.




Figure 4.337 Routing > IPv6 > Neighbor Table > Add

The following table describes the items in the previous figure.

Item	Description
Interface	The local interface on which the neighbor was discovered.
IPv6 Address	The IPv6 prefix and prefix length of the neighbor interface.
MAC Address	The MAC address associated with the neighbor interface. If the MAC address is all zeros, the entry is a Negative NDP entry. A Negative NDP entry is added to the table when the device sends a Neighbor Solicitation Request, but it has not yet been resolved. If the request is resolved and the neighbor is reachable, its valid MAC address replaces the null address. If the request times out, the entry is removed.

Item	Description
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.5.5 IPv6 Routes

4.5.5.1 IPv6 Route Table

The IPv6 Route Table page displays the entries in the IPv6 routing table, including all dynamically learned and statically configured entries. The device uses the routing table to determine how to forward IPv6 packets. A statically-configured route does not appear in the table until it is reachable.

To access this page, click **Routing > IPv6 Routes > IPv6 Route Table**.




Figure 4.338 Routing > IPv6 Routes > IPv6 Route Table

The following table describes the items in the previous figure.

Item	Description
IPv6 Prefix/Length	The IPv6 address, including the prefix and prefix length, for the destination network.
Protocol	Identifies which protocol created the route. A route can be created one of the following ways: <ul style="list-style-type: none"> ■ Dynamically learned through a supported routing protocol ■ Dynamically learned by being a directly-attached local route ■ Statically configured by an administrator
Next Hop IPv6 Address	The outgoing router IPv6 address to use when forwarding traffic to the next router (if any) in the path towards the destination. The next router is always one of the adjacent neighbors or the IPv6 address of the local interface for a directly-attached network.
Next Hop Interface	The outgoing interface to use when forwarding traffic to the destination. For a static reject route, the next hop is Null.
Best Route	Indicates whether the route is the preferred route to the network. If the field is blank, a better route to the same network exists in the IPv6 routing table.
Refresh	Click Refresh to update the screen.

4.5.5.2 IPv6 Configured Routes

Use the IPv6 Configured Routes page to configure static IPv6 global, link local, and static reject routes in the routing table. The page shows the routes that have been manually added to the routing table.

To access this page, click **Routing > IPv6 Routes > IPv6 Configured Routes**.




Figure 4.339 Routing > IPv6 Routes > IPv6 Configured Routes

The following table describes the items in the previous figure.

Item	Description
IPv6 Prefix/Length	The IPv6 address, including the prefix and prefix length, for the destination network.
Next Hop IPv6 Address	The outgoing router IPv6 address to use when forwarding traffic to the next router (if any) in the path towards the destination. The next router is always one of the adjacent neighbors or the IPv6 address of the local interface for a directly-attached network.
Next Hop Interface	The outgoing interface to use when forwarding traffic to the destination. For a static reject route, the next hop is Null. The next hop is Unresolved until the device is able to reach the interface.
Preference	The preference of the route. A lower preference value indicates a more preferred route. When the routing table has more than one route to the same network, the device selects the route with the best (lowest) route preference.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new IPv6 route. See the following procedure.
Remove	Click Remove to remove the selected entries.

To add a new IPv6 route:

Click **Routing > IPv6 Routes > IPv6 Configured Routes > Add**.

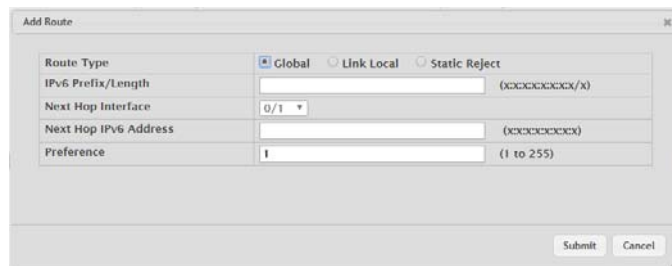


Figure 4.340 Routing > IPv6 Routes > IPv6 Configured Routes > Add

The following table describes the items in the previous figure.

Item	Description
Route Type	The type of route to configure, which is one of the following: <ul style="list-style-type: none"> ■ Global: A route with an address that is globally routable and is recognized outside of the local network. ■ Link Local: A route with an address that is allocated from part of the IPv6 unicast address space. it is not visible off the local link and is not globally unique. ■ Static Reject: A route where packets that match the route are discarded instead of forwarded. The device might send an ICMPv6 Destination Unreachable message.
IPv6 Prefix/Length	The IPv6 address, including the prefix and prefix length, for the destination network.
Next Hop Interface	The outgoing interface to use when forwarding traffic to the destination. For a static reject route, the next hop is Null. The next hop is Unresolved until the device is able to reach the interface.
Next Hop IPv6 Address	The outgoing router IPv6 address to use when forwarding traffic to the next router (if any) in the path towards the destination. The next router is always one of the adjacent neighbors or the IPv6 address of the local interface for a directly-attached network.

Item	Description
Preference	The preference of the route. A lower preference value indicates a more preferred route. When the routing table has more than one route to the same network, the device selects the route with the best (lowest) route preference.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.5.5.3 IPv6 ECMP Group

The IPv6 ECMP Groups Summary page displays all current Equal Cost Multipath (ECMP) groups in the IPv6 routing table. An ECMP group is a set of two or more next hops used in one or more routes.

To access this page, click **Routing > IPv6 Routes > IPv6 ECMP Group**.




Figure 4.341 Routing > IPv6 Routes > IPv6 ECMP Group

The following table describes the items in the previous figure.

Item	Description
ECMP Group	The ECMP group number associated with the rest of the data in the row. The device assigns an arbitrary ECMP group number from 1 to n to identify the ECMP group.
Number Of Next Hops	The number of next hops in the group.
Route(s) Count	The number of routes that use the set of next hops.
Next Hops	The IPv6 address and outgoing interface of each next hop in the group.
Refresh	Click Refresh to update the screen.

4.5.5.4 IPv6 Route Summary

The IPv6 Route Summary page displays summary information about the entries in the IPv6 routing table.

To access this page, click **Routing > IPv6 Routes > IPv6 Route Summary**.

Route Types	
Connected Routes	0
Static Routes	0
6To4 Routes	0
OSPF Routes	0
Intra Area Routes	0
Inter Area Routes	0
External Type-1 Routes	0
External Type-2 Routes	0
Total Routes	0
Route Table Counters	
Best Routes (High)	0 (0)
Alternate Routes	0
Route Adds	0
Route Deletes	0
Unresolved Route Adds	0
Invalid Route Adds	0
Failed Route Adds	0
Reserved Locals	0
Unique Next Hops (High)	0 (0)
Next Hop Groups (High)	0 (0)
ECMP Groups (High)	0 (0)
ECMP Routes	0
Truncated ECMP Routes	0
ECMP Retries	0
Number of Prefixes	

Refresh Clear Counters

Figure 4.342 Routing > IPv6 Routes > IPv6 Route Summary

The following table describes the items in the previous figure.

Item	Description
Route Types	
Connected Routes	The total number of connected routes in the IPv6 routing table.
Static Routes	The total number of static routes in the IPv6 routing table.
6To4 Routes	The total number of 6to4 routes in the IPv6 routing table. A 6to4 route allows IPv6 sites to communicate with each other over an IPv4 network by treating the wide-area IPv4 network as a unicast point-to-point link layer.
OSPF Routes	The total number of routes installed by the OSPFv3 protocol. <ul style="list-style-type: none"> ■ Intra Area Routes: The total number of intra-area routes installed by the OSPFv3 protocol. ■ Inter Area Routes: The total number of inter-area routes installed by the OSPFv3 protocol. ■ External Type-1 Routes: The total number of external type-1 routes installed by the OSPFv3 protocol. ■ External Type-2 Routes: The total number of external type-2 routes installed by the OSPFv3 protocol.
Total Routes	The total number of routes in the routing table.
Route Table Counters	
Best Routes (High)	The number of best routes currently in the routing table. This number counts only the best route to each destination.
Alternate Routes	The number of alternate routes currently in the routing table. An alternate route is a route that was not selected as the best route to its destination.
Route Adds	The number of routes that have been added to the routing table.
Route Deletes	The number of routes that have been deleted from the routing table.

Item	Description
Unresolved Route Adds	The number of route adds that failed because none of the route's next hops were on a local subnet. Note that static routes can fail to be added to the routing table at startup because the routing interfaces are not yet up. This counter gets incremented in this case. The static routes are added to the routing table when the routing interfaces come up.
Invalid Route Adds	The number of routes that failed to be added to the routing table because the route was invalid. A log message is written for each of these failures.
Failed Route Adds	The number of routes that failed to be added to the routing table because of a resource limitation in the routing table.
Reserved Locals	The number of routing table entries reserved for a local subnet on a routing interface that is down. Space for local routes is always reserved so that local routes can be installed when a routing interface bounces.
Unique Next Hops (High)	The number of distinct next hops used among all routes currently in the routing table. These include local interfaces for local routes and neighbors for indirect routes.
Next Hop Groups (High)	The current number of next hop groups in use by one or more routes. Each next hop group includes one or more next hops.
ECMP Groups (High)	The number of next hop groups with multiple next hops.
ECMP Routes	The number of routes with multiple next hops currently in the routing table.
Truncated ECMP Routes	The number of ECMP routes that are currently installed in the forwarding table with just one next hop. The forwarding table may limit the number of ECMP routes or the number of ECMP groups. When an ECMP route cannot be installed because such a limit is reached, the route is installed with a single next hop.
ECMP Retries	The number of ECMP routes that have been installed in the forwarding table after initially being installed with a single next hop.
Number of Prefixes	The unique IPv6 prefixes in the IPv6 routing table.
Refresh	Click Refresh to update the screen.
Clear Counters	Click Clear Counters to resets all IPv6 routing table event counters on this page to zero. Not that only event counters are reset; counters that report the current state of the routing table, such as the number of routes of each type, are not reset.

4.5.6 DHCPv6

4.5.6.1 Global

Use the DHCPv6 Global Configuration page to configure the global Dynamic Host Configuration Protocol for IPv6 (DHCPv6) server settings on the device. The device can act as a DHCPv6 server or DHCPv6 relay agent to help assign network configuration information to IPv6 clients.

To access this page, click **Routing > DHCPv6 > Global**.

Figure 4.343 Routing > DHCPv6 > Global

The following table describes the items in the previous figure.

Item	Description
DHCPv6 Admin Mode	The administrative mode of the DHCPv6 server.
DHCPv6 Server DUID	The DHCP Unique Identifier (DUID) of the DHCPv6 server.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.5.6.2 Pool Summary

Use the DHCPv6 Pool Summary page to view the currently configured DHCPv6 server pools and to add and remove pools. A DHCPv6 server pool is a set of network configuration information available to DHCPv6 clients that request the information.

To access this page, click **Routing > DHCPv6 > Pool Summary**.




Figure 4.344 Routing > DHCPv6 > Pool Summary

The following table describes the items in the previous figure.

Item	Description
Pool Name	The name that identifies the DHCPv6 server pool.
Delegated Prefixes	The general prefix in the pool for use in allocating and assigning addresses to hosts that may be utilizing IPv6 auto-address configuration or acting as DHCPv6 clients.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new pool. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

To add a new pool:

Click **Routing > DHCPv6 > Pool Summary > Add**.




Figure 4.345 Routing > DHCPv6 > Pool Summary > Add

The following table describes the items in the previous figure.

Item	Description
Pool Name	The name that identifies the DHCPv6 server pool.
Delegated Prefixes	The general prefix in the pool for use in allocating and assigning addresses to hosts that may be utilizing IPv6 auto-address configuration or acting as DHCPv6 clients.

Item	Description
DHCPv6 Client DUID	The DHCP Unique Identifier (DUID) of the client. The DUID is a combination of the client's hardware address and client identifier.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.5.6.3 Pool Configuration

Use the DHCPv6 Pool Configuration page to edit pool settings or to configure additional settings for existing DHCPv6 pools.

To access this page, click **Routing > DHCPv6 > Pool Configuration**.

Figure 4.346 Routing > DHCPv6 > Pool Configuration

The following table describes the items in the previous figure.

Item	Description
Pool Name	The menu includes all DHCPv6 server pools that have been configured on the device.
Delegated Prefixes	The IPv6 prefix and prefix length to assign the requesting client.
DUID	The DHCP Unique Identifier (DUID) of the client. The DUID is a combination of the client's hardware address and client identifier.
Client Name	The optional system name associated with the client.
Valid Lifetime	The maximum amount of time the requesting client is allowed to use the prefix.
Prefer Lifetime	The preferred amount of time the requesting client is allowed to use the prefix. The value of the Prefer Lifetime must be less than the value of the Valid Lifetime.
DNS Server	The IPv6 prefix of each DNS server each client in the pool can contact to perform address resolution.
Domain Name	The domain name configured for each client in the pool.
Refresh	Click Refresh to update the screen.

4.5.6.4 Interface

Use the DHCPv6 Interface Summary page to view the per-interface settings for DHCPv6. To configure the settings, select the interface to configure and click **Edit**. You are redirected to the DHCPv6 Interface Configuration page for the selected interface.

To access this page, click **Routing > DHCPv6 > Interface**.

Interface	Interface Mode	Pool Name	Relay interface	Destination IP Address	Remote ID
0/1	None	N/A	N/A	N/A	N/A
0/2	None	N/A	N/A	N/A	N/A
0/3	None	N/A	N/A	N/A	N/A
0/4	None	N/A	N/A	N/A	N/A
0/5	None	N/A	N/A	N/A	N/A
0/6	None	N/A	N/A	N/A	N/A
0/7	None	N/A	N/A	N/A	N/A
0/8	None	N/A	N/A	N/A	N/A
0/9	None	N/A	N/A	N/A	N/A
0/10	None	N/A	N/A	N/A	N/A

Figure 4.347 Routing > DHCPv6 > Interface

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row.
Interface Mode	The DHCPv6 function configured on the interface, which is one of the following: <ul style="list-style-type: none"> None: The interface is not configured as a DHCPv6 server or DHCPv6 relay agent. Server: The interface responds to requests from DHCPv6 clients. Client: The initiates requests on a link to obtain configuration parameters from one or more DHCPv6 servers. Relay: The interface acts as an intermediary to deliver DHCPv6 messages between clients and servers. The interface is on the same link as the client.
Pool Name	The name of the DHCPv6 pool the server uses to assign client information (for DHCPv6 server interface only).
Relay Interface	The interface on the device through which a DHCPv6 server is reached (for DHCPv6 relay agent interface only).
Destination IP Address	The destination IPv6 address of the DHCPv6 server to which client packets are forwarded (for DHCPv6 relay agent interface only).
Remote ID	The relay agent information option remote-ID sub-option to be added to relayed messages. This value is derived from the DHCPv6 server DUID and the relay interface number, or it can be specified as a user-defined string (for DHCPv6 relay agent interface only).
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.

4.5.6.5 Interface Configuration

Use the DHCPv6 Interface Configuration page to configure the per-interface settings for DHCPv6. With the larger address space inherent to IPv6, addresses within a network can be allocated more effectively in a hierarchical fashion. DHCPv6 introduces the notion of prefix delegation as described in RFC 3633 as a way for devices to centralize and delegate IPv6 address assignment. An interface can act as an IPv6 prefix delegation server that defines one or more general prefixes to delegate to a device lower in the hierarchy acting as a prefix delegation client. The device with an interface configured as a prefix delegation client can then allocate more specific addresses within the given general prefix range to assign to its local router interfaces. This device can, in turn, use the given general prefix in allocating and assigning addresses to host machines that may be utilizing IPv6 auto-address configuration or acting as DHCPv6 clients. An interface can also be configured as a DHCPv6 relay agent that acts as an intermediary to deliver DHCPv6 messages between clients and servers and is on the same link as the client. The DHCPv6 interface modes are mutually exclusive. The fields that can be configured on this page depend on the selected mode for the interface.

To access this page, click **Routing > DHCPv6 > Interface Configuration**.

Figure 4.348 Routing > DHCPv6 > Interface Configuration

The following table describes the items in the previous figure.

Item	Description
Interface	Select the interface with the information to view or configure.
Interface Mode	The DHCPv6 function configured on the interface, which is one of the following: <ul style="list-style-type: none"> None: The interface is not configured as a DHCPv6 server or DHCPv6 client or DHCPv6 relay agent. Server: The interface responds to requests from DHCPv6 clients. Client: The interface initiates requests on a link to obtain configuration parameters from one or more DHCPv6 servers. Relay: The interface acts as an intermediary to deliver DHCPv6 messages between clients and servers. The interface is on the same link as the client.
Server Options	
Pool Name	The name of the DHCPv6 pool the server can use to assign client information.

Item	Description
Rapid Commit	The mode of the rapid commit message exchange on the DHCPv6 server interface. The DHCPv6 client can obtain configuration parameters from a server either through a rapid two-message exchange (solicit, reply) or through a four-message exchange (solicit, advertise, request, and reply). When this option is disabled on either the client or server, the four-message exchange is used. When this option is enabled on both the client and the server, the two-message exchange is used.
Preference	The preference value to include in DHCPv6 Advertise messages. If a DHCPv6 client receives Advertise messages from multiple DHCPv6 servers, it responds to the server with the highest preference value.
Client Options	
Prefix Delegation Client	When enabled, the interface can receive a general prefix for assignment to local router interfaces.
Rapid Commit	The mode of the rapid commit message exchange on the DHCPv6 client interface. The DHCPv6 client can obtain configuration parameters from a server either through a rapid two-message exchange (solicit, reply) or through a four-message exchange (solicit, advertise, request, and reply). When this option is disabled on either the client or server, the four-message exchange is used. When this option is enabled on both the client and the server, the two-message exchange is used.
Relay Options	
Relay Interface	The interface on the device through which a DHCPv6 server is reached.
Destination IP Address	The destination IPv6 address of the DHCPv6 server to which client packets are forwarded.
Remote ID	The relay agent information option remote-ID sub-option to be added to relayed messages. This value is derived from the DHCPv6 server DUID and the relay interface number, or it can be specified as a user-defined string.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.5.6.6 Bindings

Use the DHCPv6 Binding Summary page to view entries in the DHCP Bindings table. After a client acquires IPv6 configuration information from the DHCPv6 server, the server adds an entry to its database. The entry is called a binding.

To access this page, click **Routing > DHCPv6 > Bindings**.



Figure 4.349 Routing > DHCPv6 > Bindings

The following table describes the items in the previous figure.

Item	Description
Client IP Address	The IPv6 address associated with the client.
Client Interface	The interface number where the client binding occurred.

Item	Description
DHCPv6 Client DUID	The DHCP Unique Identifier (DUID) of the client. The DUID is a combination of the client's hardware address and client identifier.
IPv6 Prefix	The type of prefix associated with this binding.
Expiry Time	The number of seconds until the prefix associated with a binding expires.
Valid Lifetime	The maximum amount of time the client is allowed to use the prefix.
Prefer Lifetime	The preferred amount of time the client is allowed to use the prefix.
Refresh	Click Refresh to update the screen.
Clear Entries	Click Clear Entries to remove an entry from the table.

4.5.6.7 Statistics

The DHCPv6 Statistics page displays the DHCPv6 server statistics for the device, including information about the DHCPv6 messages sent, received, and discarded globally and on each interface. The values on this page indicate the various counts that have accumulated since they were last cleared.

To access this page, click **Routing > DHCPv6 > Statistics**.

Interface	Total DHCPv6 Packets Received	DHCPv6 Request Packets Received	Received DHCPv6 Packets Discarded	Total DHCPv6 Packets Sent	DHCPv6 Reply Packets Transmitted
All	0	0	0	0	0
0/1	0	0	0	0	0
0/2	0	0	0	0	0
0/3	0	0	0	0	0
0/4	0	0	0	0	0
0/5	0	0	0	0	0
0/6	0	0	0	0	0
0/7	0	0	0	0	0
0/8	0	0	0	0	0
0/9	0	0	0	0	0

Figure 4.350 Routing > DHCPv6 > Statistics

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. The row at the top of the table (All) contains cumulative statistics for all interfaces.
Total DHCPv6 Packets Received	The number of DHCPv6 messages received on the interface. The DHCPv6 messages sent from a DHCPv6 client to a DHCPv6 server include Solicit, Request, Confirm, Renew, Rebind, Release, Decline, and Information-Request messages. Additionally, a DHCPv6 relay agent can forward Relay-Forward messages to a DHCPv6 server.
DHCPv6 Request Packets Received	The number of DHCPv6 Request messages received on the interface. DHCPv6 Request messages are sent by a client to request IPv6 configuration information from the server.
Received DHCPv6 Packets Discarded	The number of DHCPv6 messages received on the interface that were discarded due to errors or because they were invalid.
Total DHCPv6 Packets Sent	The number of DHCPv6 messages sent by the interface. The DHCPv6 messages sent from a DHCPv6 server to a DHCPv6 client include Advertise, Reply, Reconfigure, and Relay-Reply messages.
DHCPv6 Reply Packets Transmitted	The number of DHCPv6 Reply messages sent from the interface to a DHCPv6 client in response to a Solicit, Request, Renew, Rebind, Information-Request, Confirm, Release, or Decline message.
Refresh	Click Refresh to update the screen.
Details	Click Details to open the Details window.

Item	Description
Clear	Click Clear to reset the DHCPv6 counters for one or more interfaces for the selected entries.

4.6 Security

4.6.1 Port Access Control

In port-based authentication mode, when 802.1x is enabled globally and on the port, successful authentication of any one supplicant attached to the port results in all users being able to use the port without restrictions. At any given time, only one supplicant is allowed to attempt authentication on a port in this mode. Ports in this mode are under bidirectional control. This is the default authentication mode.

The 802.1X network has three components:

- Authenticators: Specifies the port that is authenticated before permitting system access.
- Supplicants: Specifies host connected to the authenticated port requesting access to the system services.
- Authentication Server: Specifies the external server, for example, the RADIUS server that performs the authentication on behalf of the authenticator, and indicates whether the user is authorized to access system services.

The Port Access Control folder contains links to the following pages that allow you to view and configure 802.1X features on the system.

4.6.1.1 Configuration

Use the Port Access Control Configuration page to enable or disable port access control on the system.

To access this page, click **Security > Port Access Control > Configuration**.

Admin Mode	<input checked="" type="radio"/> Disable <input type="radio"/> Enable
VLAN Assignment Mode	<input checked="" type="radio"/> Disable <input type="radio"/> Enable
Dynamic VLAN Creation Mode	<input checked="" type="radio"/> Disable <input type="radio"/> Enable
Monitor Mode	<input checked="" type="radio"/> Disable <input type="radio"/> Enable
EAPOL Flood Mode	Disabled ✓

Submit Refresh Cancel

Figure 4.351 Security > Port Access Control > Configuration

The following table describes the items in the previous figure.

Item	Description
Admin Mode	The administrative mode of port-based authentication on the device.
VLAN Assignment Mode	The administrative mode of RADIUS-based VLAN assignment on the device. When enabled, this feature allows a port to be placed into a particular VLAN based on the result of the authentication or type of 802.1X authentication a client uses when it accesses the device. The authentication server can provide information to the device about which VLAN to assign the supplicant.
Dynamic VLAN Creation Mode	The administrative mode of dynamic VLAN creation on the device. If RADIUS-assigned VLANs are enabled, the RADIUS server is expected to include the VLAN ID in the 802.1X tunnel attributes of its response message to the device. If dynamic VLAN creation is enabled on the device and the RADIUS-assigned VLAN does not exist, then the assigned VLAN is dynamically created. This implies that the client can connect from any port and can get assigned to the appropriate VLAN. This feature gives flexibility for clients to move around the network without much additional configuration required.

Item	Description
Monitor Mode	The administrative mode of the Monitor Mode feature on the device. Monitor mode is a special mode that can be enabled in conjunction with port-based access control. Monitor mode provides a way for network administrators to identify possible issues with the port-based access control configuration on the device without affecting the network access to the users of the device. It allows network access even in cases where there is a failure to authenticate, but it logs the results of the authentication process for diagnostic purposes. If the device fails to authenticate a client for any reason (for example, RADIUS access reject from the RADIUS server, RADIUS timeout, or the client itself is 802.1X unaware), the client is authenticated and is undisturbed by the failure condition(s). The reasons for failure are logged and buffered into the local logging database for tracking purposes.
EAPOL Flood Mode	The administrative mode of the Extensible Authentication Protocol (EAP) over LAN (EAPOL) flood support on the device. EAPOL Flood Mode can be enabled when Admin Mode and Monitor Mode are disabled.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.6.1.2 Port Summary

Use the Port Access Control Port Summary page to view summary information about the port-based authentication settings for each port.



To access this page, click **Security > Port Access Control > Port Summary**.

Interface	PAE Capabilities	Control Mode	Operating Control Mode	PAE State	Backend State
0/1	Authenticator	Auto	N/A	Initialize	Initialize
0/2	Authenticator	Auto	N/A	Initialize	Initialize
0/3	Authenticator	Auto	N/A	Initialize	Initialize
0/4	Authenticator	Auto	N/A	Initialize	Initialize
0/5	Authenticator	Auto	N/A	Initialize	Initialize
0/6	Authenticator	Auto	N/A	Initialize	Initialize
0/7	Authenticator	Auto	N/A	Initialize	Initialize
0/8	Authenticator	Auto	N/A	Initialize	Initialize
0/9	Authenticator	Auto	N/A	Initialize	Initialize
0/10	Authenticator	Auto	N/A	Initialize	Initialize

Figure 4.352 Security > Port Access Control > Port Summary

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row.
PAE Capabilities	The Port Access Entity (PAE) role, which is one of the following: <ul style="list-style-type: none"> Authenticator: The port enforces authentication and passes authentication information from a remote supplicant (similar to a client or host) to the authentication server. If the server successfully authenticates the supplicant, the port allows access. Supplicant: The port must be granted permission by the authentication server before it can access the remote authenticator port.

Item	Description
Control Mode	<p>The port-based access control mode configured on the port, which is one of the following:</p> <ul style="list-style-type: none"> ■ Auto: The port is unauthorized until a successful authentication exchange has taken place. ■ Force Unauthorized: The port ignores supplicant authentication attempts and does not provide authentication services to the client. ■ Force Authorized: The port sends and receives normal traffic without client port-based authentication.
Operating Control Mode	<p>The control mode under which the port is actually operating, which is one of the following:</p> <ul style="list-style-type: none"> ■ Auto ■ Force Unauthorized ■ Force Authorized ■ N/A <p>If the mode is N/A, port-based access control is not applicable to the port. If the port is in detached state it cannot participate in port access control. Additionally, if port-based access control is globally disabled, the status for all ports is N/A.</p>
PAE State	<p>The current state of the authenticator PAE state machine, which is the 802.1X process that controls access to the port. The state can be one of the following:</p> <ul style="list-style-type: none"> ■ Initialize ■ Disconnected ■ Connecting ■ Authenticating ■ Authenticated ■ Aborting ■ Held ■ ForceAuthorized ■ ForceUnauthorized
Backend State	<p>The current state of the backend authentication state machine, which is the 802.1X process that controls the interaction between the 802.1X client on the local system and the remote authentication server. The state can be one of the following:</p> <ul style="list-style-type: none"> ■ Request ■ Response ■ Success ■ Fail ■ Timeout ■ Initialize ■ Idle
	<p>Click the Initialize button to reset the 802.1X state machine on the associated interface to the initialization state. Traffic sent to and from the port is blocked during the authentication process. This icon can be clicked only when the port is an authenticator and the operating control mode is Auto.</p>
	<p>Click the Re-Authenticate button to force the associated interface to restart the authentication process.</p>
Refresh	<p>Click Refresh to update the screen.</p>
Edit	<p>Click Edit to edit the selected entries.</p>
Details	<p>Click Details to open a window and display additional information.</p>

4.6.1.3 Port Configuration

Use the Port Access Control Port Configuration page to enable and configure port access control on one or more ports.

To access this page, click **Security > Port Access Control > Port Configuration**.

The screenshot shows the 'Port Configuration' page for interface 0/1. The 'PAE Capabilities' is set to 'Authenticator'. The 'Authenticator Options' section includes: Control Mode (Auto), Quiet Period (60s), Transmit Period (30s), Guest VLAN ID (empty), Guest VLAN Period (90s), Unauthenticated VLAN ID (empty), Supplicant Timeout (30s), Server Timeout (30s), Maximum Requests (2), MAB Mode (unchecked), Re-Authentication Period (Disabled), and Maximum Users (48). The 'Supplicant Options' section includes: Control Mode (Auto), User Name (None), Authentication Period (30s), Start Period (30s), Held Period (60s), and Maximum Start Messages (3). Buttons for 'Submit', 'Refresh', and 'Cancel' are at the bottom.

Figure 4.353 Security > Port Access Control > Port Configuration

The following table describes the items in the previous figure.

Item	Description
Interface	The interface with the settings to view or configure. If you have been redirected to this page, this field is read-only and displays the interface that was selected on the Port Access Control Port Summary page.
PAE Capabilities	<p>The Port Access Entity (PAE) role, which is one of the following:</p> <ul style="list-style-type: none"> Authenticator: The port enforces authentication and passes authentication information from a remote supplicant (client or host) to the authentication server. If the server successfully authenticates the supplicant, the port allows access. Supplicant: The port is connected to an authenticator port and must be granted permission by the authentication server before it can send and receive traffic through the remote port. <p>To change the PAE capabilities of a port, click the Edit icon associated with the field and select the desired setting from the menu in the Set PAE Capabilities window.</p>
Authenticator Options	
Control Mode	<p>The port-based access control mode on the port, which is one of the following:</p> <ul style="list-style-type: none"> Auto: The port is unauthorized until a successful authentication exchange has taken place. Force Unauthorized: The port ignores supplicant authentication attempts and does not provide authentication services to the client. Force Authorized: The port sends and receives normal traffic without client port-based authentication. MAC-Based: This mode allows multiple supplicants connected to the same port to each authenticate individually. Each host connected to the port must authenticate separately in order to gain access to the network. The hosts are distinguished by their MAC addresses.

Item	Description
Quiet Period (Seconds)	The number of seconds that the port remains in the quiet state following a failed authentication exchange.
Transmit Period (Seconds)	The value, in seconds, of the timer used by the authenticator state machine on the port to determine when to send an EAPOL EAP Request/Identity frame to the supplicant.
Guest VLAN ID	The VLAN ID for the guest VLAN. The guest VLAN allows the port to provide a distinguished service to unauthenticated users. This feature provides a mechanism to allow users access to hosts on the guest VLAN. To set the guest VLAN ID, click the Edit icon associated with the field and specify the ID value in the available field. To reset the guest VLAN ID to the default value, click the Reset icon associated with the field and confirm the action.
Guest VLAN Period (Seconds)	The value, in seconds, of the timer used for guest VLAN authentication.
Unauthenticated VLAN ID	The VLAN ID of the unauthenticated VLAN. Hosts that fail the authentication might be denied access to the network or placed on a VLAN created for unauthenticated clients. This VLAN might be configured with limited network access. To set the unauthenticated VLAN ID, click the Edit icon associated with the field and specify the ID value in the available field. To reset the unauthenticated VLAN ID to the default value, click the Reset icon associated with the field and confirm the action.
Supplicant Timeout (Seconds)	The amount of time that the port waits for a response before retransmitting an EAP request frame to the client.
Server Timeout (Seconds)	The amount of time the port waits for a response from the authentication server.
Maximum Requests	The maximum number of times that the port sends an EAP request frame (assuming that no response is received) to the client before restarting the authentication process.
MAB Mode	The MAC-based Authentication Bypass (MAB) mode on the port, which can be enabled or disabled.
Re-Authentication Period (Seconds)	The amount of time that clients can be connected to the port without being reauthenticated. If this field is disabled, connected clients are not forced to reauthenticate periodically. To change the value, click the Edit icon associated with the field and specify a value in the available field. To reset the reauthentication period to the default value, click the Reset icon associated with the field and confirm the action.
Maximum Users	The maximum number of clients supported on the port if the Control Mode on the port is MAC-based 802.1X authentication.
Supplicant Options	
Control Mode	The port-based access control mode on the port, which is one of the following: <ul style="list-style-type: none"> ■ Auto: The port is in an unauthorized state until a successful authentication exchange has taken place between the supplicant port, the authenticator port, and the authentication server. ■ Force Unauthorized: The port is placed into an unauthorized state and is automatically denied system access. ■ Force Authorized: The port is placed into an authorized state and does not require client port-based authentication to be able to send and receive traffic.
User Name	The name the port uses to identify itself as a supplicant to the authenticator port. The menu includes the users that are configured for system management. When authenticating, the supplicant provides the password associated with the selected User Name.

Item	Description
Authentication Period (Seconds)	The amount of time the supplicant port waits to receive a challenge from the authentication server. If the configured Authentication Period expires, the supplicant retransmits the authentication request until it is authenticated or has sent the number of messages configured in the Maximum Start Messages field.
Start Period (Seconds)	The amount of time the supplicant port waits for a response from the authenticator port after sending a Start packet. If no response is received, the supplicant retransmits the Start packet.
Held Period (Seconds)	The amount of time the supplicant port waits before contacting the authenticator port after an active 802.1X session fails.
Maximum Start Messages	The maximum number of Start packets the supplicant port sends to the authenticator port without receiving a response before it considers the authenticator to be 802.1X-unaware.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.6.1.4 Port Details

Use the Port Access Control Port Details page to view 802.1X information for a specific port.

To access this page, click **Security > Port Access Control > Port Details**.

Authenticator Options	
Control Mode	Auto
Quiet Period (Seconds)	60
Transmit Period (Seconds)	30
Guest VLAN ID	0
Guest VLAN Period (Seconds)	90
Unauthenticated VLAN ID	0
Supplicant Timeout (Seconds)	30
Server Timeout (Seconds)	30
Maximum Requests	2
Configured MAB Mode	Disabled
Operational MAB Mode	Disabled
Re-Authentication Period (Seconds)	Disabled
Maximum Users	48

Figure 4.354 Security > Port Access Control > Port Details

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data on the page.
PAE Capabilities	The Port Access Entity (PAE) role, which is one of the following: <ul style="list-style-type: none"> ■ Authenticator: The port enforces authentication and passes authentication information from a remote supplicant (client or host) to the authentication server. If the server successfully authenticates the supplicant, the port allows access. ■ Supplicant: The port is connected to an authenticator port and must be granted permission by the authentication server before it can send and receive traffic through the remote port.

Item	Description
Authenticator Options	
Control Mode	The port-based access control mode on the port, which is one of the following: <ul style="list-style-type: none"> ■ Auto: The port is unauthorized until a successful authentication exchange has taken place. ■ Force Unauthorized: The port ignores supplicant authentication attempts and does not provide authentication services to the client. ■ Force Authorized: The port sends and receives normal traffic without client port-based authentication.
Quiet Period (Seconds)	The number of seconds that the port remains in the quiet state following a failed authentication exchange.
Transmit Period (Seconds)	The value, in seconds, of the timer used by the authenticator state machine on the port to determine when to send an EAPOL EAP Request/Identity frame to the supplicant.
Guest VLAN ID	The VLAN ID for the guest VLAN. The guest VLAN allows the port to provide a distinguished service to unauthenticated users. This feature provides a mechanism to allow users access to hosts on the guest VLAN.
Guest VLAN Period (Seconds)	The value, in seconds, of the timer used for guest VLAN authentication.
Unauthenticated VLAN ID	The VLAN ID of the unauthenticated VLAN. Hosts that fail the authentication might be denied access to the network or placed on a VLAN created for unauthenticated clients. This VLAN might be configured with limited network access.
Supplicant Timeout (Seconds)	The amount of time that the port waits for a response before retransmitting an EAP request frame to the client.
Server Timeout (Seconds)	The amount of time the port waits for a response from the authentication server.
Maximum Requests	The maximum number of times that the port sends an EAP request frame (assuming that no response is received) to the client before restarting the authentication process.
Configured MAB Mode	The configured MAC-based Authentication Bypass (MAB) mode on the port.
Operational MAB Mode	The operational MAC-based Authentication Bypass (MAB) mode on the port.
Re-Authentication Period (Seconds)	The amount of time that clients can be connected to the port without being reauthenticated. If this field is disabled, connected clients are not forced to reauthenticate periodically.
Maximum Users	The maximum number of clients supported on the port if the Control Mode on the port is MAC-based 802.1X authentication.
Refresh	Click Refresh to update the screen.

4.6.1.5 Statistics

Use the Port Access Control Statistics page to view information about the Extensible Authentication Protocol over LAN (EAPOL) frames and EAP messages sent and received by the local interfaces.

To access this page, click **Security > Port Access Control > Statistics**.

Interface	PAE Capabilities	EAPOL Frames Received	EAPOL Frames Transmitted	Last EAPOL Frame Version	Last EAPOL Frame Source
0/1	Authenticator	0	0	0	00:00:00:00:00:00
0/2	Authenticator	0	0	0	00:00:00:00:00:00
0/3	Authenticator	0	0	0	00:00:00:00:00:00
0/4	Authenticator	0	0	0	00:00:00:00:00:00
0/5	Authenticator	0	0	0	00:00:00:00:00:00
0/6	Authenticator	0	0	0	00:00:00:00:00:00
0/7	Authenticator	0	0	0	00:00:00:00:00:00
0/8	Authenticator	0	0	0	00:00:00:00:00:00
0/9	Authenticator	0	0	0	00:00:00:00:00:00
0/10	Authenticator	0	0	0	00:00:00:00:00:00

Figure 4.355 Security > Port Access Control > Statistics

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. When viewing detailed information for an interface, this field identifies the interface being viewed.
PAE Capabilities	The Port Access Entity (PAE) role, which is one of the following: <ul style="list-style-type: none"> Authenticator: The port enforces authentication and passes authentication information from a remote supplicant (similar to a client or host) to the authentication server. If the server successfully authenticates the supplicant, the port allows access. Supplicant: The port must be granted permission by the authentication server before it can access the remote authenticator port.
EAPOL Frames Received	The total number of valid EAPOL frames received on the interface.
EAPOL Frames Transmitted	The total number of EAPOL frames sent by the interface.
Last EAPOL Frame Version	The protocol version number attached to the most recently received EAPOL frame.
Last EAPOL Frame Source	The source MAC address attached to the most recently received EAPOL frame.
Refresh	Click Refresh to update the screen.
Details	Click Details to open a window and display additional information.
Clear	Click Clear to reset all statistics counters to zero.

4.6.1.6 Client Summary

The Port Access Control Client Summary page displays information about supplicant devices that are connected to the local authenticator ports. If there are no active 802.1X sessions, the table is empty.

To access this page, click **Security > Port Access Control > Client Summary**.




Figure 4.356 Security > Port Access Control > Client Summary

The following table describes the items in the previous figure.

Item	Description
Interface	The local interface associated with the rest of the data in the row. When viewing detailed information for an interface, this field identifies the interface being viewed.
Logical Interface	The logical port number associated with the supplicant that is connected to the port.
User Name	The name the client uses to identify itself as a supplicant to the authentication server.
Supplicant MAC Address	The MAC address of the supplicant that is connected to the port.
Session Time	The amount of time that has passed since the connected supplicant was granted access to the network through the authenticator port.
Filter ID	The policy filter ID assigned by the authenticator to the supplicant device.
VLAN ID	The ID of the VLAN the supplicant was placed in as a result of the authentication process.
Refresh	Click Refresh to update the screen.
Details	Click Details to open a window and display additional information.

4.6.1.7 Privileges Summary

Use the Port Access Control Privileges Summary page to grant or deny port access to users configured on the system.

To access this page, click **Security > Port Access Control > Privileges Summary**.

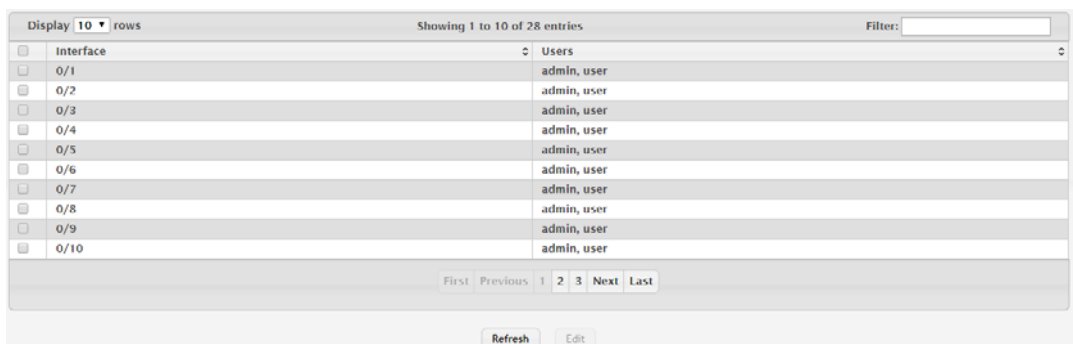


Figure 4.357 Security > Port Access Control > Privileges Summary

The following table describes the items in the previous figure.

Item	Description
Interface	The local interface associated with the rest of the data in the row. When configuring access information for one or more interfaces, this field identifies each interface being configured.
Users	The users that are allowed access to the system through the associated port. When configuring user access for a port, the Available Users field lists the users configured on the system that are denied access to the port. The users in the Selected Users field are allowed access. To move a user from one field to the other, click the user to move (or CTL + click to select multiple users) and click the appropriate arrow.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.

4.6.1.8 History Log Summary

Use the Port Access Control History Log Summary page to grant or deny port access to users configured on the system.

To access this page, click **Security > Port Access Control > History Log Summary**.




Figure 4.358 Security > Port Access Control > History Log Summary

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. Only interfaces that have entries in the log history are listed.
Time Stamp	The absolute time when the authentication event took place.
VLAN Assigned	The ID of the VLAN the supplicant was placed in as a result of the authentication process.
VLAN Assigned Reason	The reason why the authenticator placed the supplicant in the VLAN. Possible values are: <ul style="list-style-type: none"> ■ RADIUS ■ Unauth ■ Default ■ Not Assigned
Supp MAC Address	The MAC address of the supplicant that is connected to the port.
Filter Name	The policy filter ID assigned by the authenticator to the supplicant device.
Auth Status	The authentication status of the client or port.
Reason	The reason for the successful or unsuccessful authentication.
Refresh	Click Refresh to update the screen.
Clear History	Click Clear History to clear the history logs.

4.6.2 RADIUS

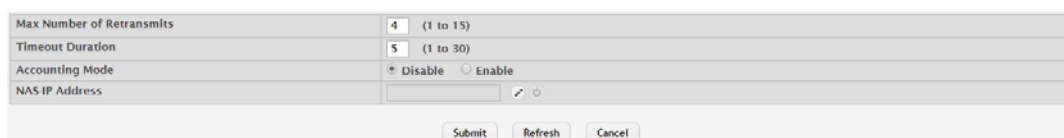
Remote Authorization Dial-In User Service (RADIUS) servers provide additional security for networks. The RADIUS server maintains a user database, which contains per-user authentication information. RADIUS servers provide a centralized authentication method for:

- Telnet Access
- Web Access
- Console to Switch Access
- Port Access Control (802.1X)

4.6.2.1 Configuration

Use the RADIUS Configuration page to view and configure various settings for the RADIUS servers configured on the system.

To access this page, click **Security > RADIUS > Configuration**.





The screenshot shows a configuration form with the following fields and values:

Max Number of Retransmits	4 (1 to 15)
Timeout Duration	5 (1 to 30)
Accounting Mode	<input checked="" type="radio"/> Disable <input type="radio"/> Enable
NAS IP Address	<input type="text"/>

Buttons: Submit, Refresh, Cancel

Figure 4.359 Security > RADIUS > Configuration

The following table describes the items in the previous figure.

Item	Description
Max Number of Retransmits	The maximum number of times the RADIUS client on the device will retransmit a request packet to a configured RADIUS server after a response is not received. If multiple RADIUS servers are configured, the max retransmit value will be exhausted on the first server before the next server is attempted. A retransmit will not occur until the configured timeout value on that server has passed without a response from the RADIUS server. Therefore, the maximum delay in receiving a response from the RADIUS server equals the sum of (retransmit - timeout) for all configured servers. If the RADIUS request was generated by a user login attempt, all user interfaces will be blocked until the RADIUS application returns a response.
Timeout Duration	The number of seconds the RADIUS client waits for a response from the RADIUS server. Consideration to maximum delay time should be given when configuring RADIUS timeout and RADIUS max retransmit values.
Accounting Mode	Specifies whether the RADIUS accounting mode on the device is enabled or disabled.
NAS-IP Address	The network access server (NAS) IP address for the RADIUS server. To specify an address, click  button and enter the IP address of the NAS in the available field. The address should be unique to the NAS within the scope of the RADIUS server. The NAS IP address is used only in Access-Request packets. To reset the NAS IP address to the default value, click  button and confirm the action.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.6.2.2 Named Server

The RADIUS Named Server Status page shows summary information about the RADIUS servers configured on the system.

To access this page, click **Security > RADIUS > Named Server**.




Figure 4.360 Security > RADIUS > Named Server

The following table describes the items in the previous figure.

Item	Description
Current	Indicates whether the RADIUS server is the current server (True) or a backup server (False) within its group. If more than one RADIUS server is configured with the same Server Name, the device selects one of the servers to be the current server in the named server group. When the device sends a RADIUS request to the named server, the request is directed to the server selected as the current server. Initially the primary server is selected as the current server. If the primary server fails, one of the other servers becomes the current server. If no server is configured as the primary server, the current server is the RADIUS server that is added to the group first.
IP Address/Host Name	The IP address or host name of the RADIUS server. Host names must be resolvable by DNS and are composed of a series of labels separated by dots.
Server Name	The name of the RADIUS server. RADIUS authentication servers that are configured with the same name are members of the same named RADIUS server group. RADIUS servers in the same group serve as backups for each other.
Port Number	The UDP port on the RADIUS authentication server to which the local RADIUS client sends request packets.
Server Type	Indicates whether the server is the Primary or a Secondary RADIUS authentication server. When multiple RADIUS servers have the same Server Name value, the RADIUS client attempts to use the primary server first. If the primary server does not respond, the RADIUS client attempts to use one of the backup servers within the same named server group.
Secret Configured	Indicates whether the shared secret for this server has been configured.
Message Authenticator	Indicates whether the RADIUS server requires the Message Authenticator attribute to be present. The Message Authenticator adds protection to RADIUS messages by using an MD5 hash to encrypt each message. The shared secret is used as the key, and if the message fails to be verified by the RADIUS server, it is discarded.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new RADIUS authentication server. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

To add a new RADIUS authentication server:
 Click **Security > RADIUS > Named Server > Add**.

Figure 4.361 Security > RADIUS > Named Server > Add

The following table describes the items in the previous figure.

Item	Description
IP Address/Host Name	The IP address or host name of the RADIUS server. Host names must be resolvable by DNS and are composed of a series of labels separated by dots.
Server Name	The name of the RADIUS server. RADIUS authentication servers that are configured with the same name are members of the same named RADIUS server group. RADIUS servers in the same group serve as backups for each other.
Port Number	The UDP port on the RADIUS authentication server to which the local RADIUS client sends request packets.
Secret	The shared secret text string used for authenticating and encrypting all RADIUS communications between the RADIUS client on the device and the RADIUS server. The secret specified in this field must match the shared secret configured on the RADIUS server.
Server Type	Indicates whether the server is the Primary or a Secondary RADIUS authentication server. When multiple RADIUS servers have the same Server Name value, the RADIUS client attempts to use the primary server first. If the primary server does not respond, the RADIUS client attempts to use one of the backup servers within the same named server group.
Message Authenticator	Indicates whether the RADIUS server requires the Message Authenticator attribute to be present. The Message Authenticator adds protection to RADIUS messages by using an MD5 hash to encrypt each message. The shared secret is used as the key, and if the message fails to be verified by the RADIUS server, it is discarded.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.6.2.3 Statistics

Use the RADIUS Server Statistics page to view statistical information for each RADIUS server configured on the system.

To access this page, click **Security > RADIUS > Statistics**.

Figure 4.362 Security > RADIUS > Statistics

The following table describes the items in the previous figure.

Item	Description
IP Address/Host Name	The IP address or host name of the RADIUS server associated with the rest of the data in the row. When viewing the detailed statistics for a RADIUS server, this field identifies the RADIUS server.
Round Trip Time	The time interval, in hundredths of a second, between the most recent Access-Reply/Access-Challenge and the Access-Request that matched it from the RADIUS authentication server.
Access Requests	The number of RADIUS Access-Request packets sent to the server. This number does not include retransmissions.
Access Rejects	The number of RADIUS Access-Reject packets, including both valid and invalid packets, that were received from the server.
Pending Requests	The number of RADIUS Access-Request packets destined for the server that have not yet timed out or received a response.
Timeouts	The number of times a response was not received from the server within the configured timeout value.
Packets Dropped	The number of RADIUS packets received from the server on the authentication port and dropped for some other reason.
Refresh	Click Refresh to update the screen.
Details	Click Details to open a window and display additional information.

4.6.2.4 Accounting Server

The RADIUS Accounting Server Status page shows summary information about the accounting servers configured on the system.

To access this page, click **Security > RADIUS > Accounting Server**.




Figure 4.363 Security > RADIUS > Accounting Server

The following table describes the items in the previous figure.

Item	Description
IP Address/Host Name	The IP address or host name of the RADIUS accounting server. Host names must be resolvable by DNS and are composed of a series of labels separated by dots.
Server Name	The name of the RADIUS accounting server. The server name must be unique among all configured RADIUS accounting servers.
Port Number	The UDP port on the RADIUS accounting server to which the local RADIUS client sends request packets.
Secret Configured	Indicates whether the shared secret for this server has been configured.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new RADIUS accounting server. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

To add a new RADIUS accounting server:

Click **Security > RADIUS > Accounting Server > Add**.

Figure 4.364 Security > RADIUS > Accounting Server > Add

The following table describes the items in the previous figure.

Item	Description
IP Address/Host Name	The IP address or host name of the RADIUS accounting server. Host names must be resolvable by DNS and are composed of a series of labels separated by dots.
Server Name	The name of the RADIUS accounting server. The server name must be unique among all configured RADIUS accounting servers.
Port Number	The UDP port on the RADIUS accounting server to which the local RADIUS client sends request packets.
Secret	The shared secret text string used for authenticating and encrypting all RADIUS communications between the RADIUS client on the device and the RADIUS accounting server. The secret specified in this field must match the shared secret configured on the RADIUS accounting server.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.6.2.5 Accounting Statistics

Use the RADIUS Accounting Server Statistics page to view statistical information for each RADIUS server configured on the system.

To access this page, click **Security > RADIUS > Accounting Statistics**.

Figure 4.365 Security > RADIUS > Accounting Statistics

The following table describes the items in the previous figure.

Item	Description
IP Address/Host Name	The IP address or host name of the RADIUS accounting server associated with the rest of the data in the row. When viewing the detailed statistics for a RADIUS accounting server, this field identifies the server.
Round Trip Time	The time interval, in hundredths of a second, between the most recent Accounting-Response and the Accounting-Request that matched it from the RADIUS accounting server.
Accounting Requests	The number of RADIUS Accounting-Request packets sent to the server. This number does not include retransmissions.
Pending Requests	The number of RADIUS Accounting-Request packets destined for the server that have not yet timed out or received a response.

Item	Description
Timeouts	The number of times a response was not received from the server within the configured timeout value.
Packets Dropped	The number of RADIUS packets received from the server on the accounting port and dropped for some other reason.
Refresh	Click Refresh to update the screen.
Details	Click Details to open a window and display additional information.

4.6.2.6 Clear Statistics

Use the RADIUS Clear Statistics page to reset all RADIUS authentication and accounting statistics to zero.

To access this page, click **Security > RADIUS > Clear Statistics**.




Figure 4.366 Security > RADIUS > Clear Statistics

The following table describes the items in the previous figure.

Item	Description
Reset	Click Reset to clear all RADIUS authentication and RAIDUS accounting server statistics.

4.6.2.7 Source Interface Configuration

Use the RADIUS Source Interface Configuration page to specify the physical or logical interface to use as the RADIUS client source interface. When an IP address is configured on the source interface, this address is used for all RADIUS communications between the local RADIUS client and the remote RADIUS server. The IP address of the designated source interface is used in the IP header of RADIUS management protocol packets. This allows security devices, such as firewalls, to identify all source packets coming from a specific device.

To access this page, click **Security > RADIUS > Source Interface Configuration**.

Figure 4.367 Security > RADIUS > Source Interface Configuration

The following table describes the items in the previous figure.

Item	Description
Type	The type of interface to use as the source interface: <ul style="list-style-type: none"> ■ None: The primary IP address of the originating (outbound) interface is used as the source address. ■ Interface: The primary IP address of a physical port is used as the source address. ■ VLAN: The primary IP address of a VLAN routing interface is used as the source address. ■ Loopback: The primary IP address of the loopback interface is used as the source address. A loopback is always reachable, as long as any routing interface is up. ■ Network: The network source IP is used as the source address. ■ Service Port: The management port source IP is used as the source address.
Interface	When the selected Type is Interface, select the physical port to use as the source interface.
VLAN ID	When the selected Type is VLAN, select the VLAN to use as the source interface. The menu contains only the VLAN IDs for VLAN routing interfaces.
Loopback Interface	When the selected Type is Loopback, select the loopback interface to use as the source interface.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.6.3 TACACS+

4.6.3.1 Configuration

Use the TACACS+ Configuration page to setup accounting information and administration control over authentication and authorization between the TACACS+ server and the device.

To access this page, click **Security > TACACS+ > Configuration**.

The screenshot shows a configuration form with two main input fields. The first field is labeled 'Key String' and has a text input area. The second field is labeled 'Connection Timeout' and has a numeric input field containing the value '5', with a note '(1 to 30 secs)' to its right. Below the input fields are three buttons: 'Submit', 'Refresh', and 'Cancel'.

Figure 4.368 Security > TACACS+ > Configuration

The following table describes the items in the previous figure.

Item	Description
Key String	Specifies the authentication and encryption key for TACACS+ communications between the device and the TACACS+ server. The key must match the key configured on the TACACS+ server.
Connection Timeout	The maximum number of seconds allowed to establish a TCP connection between the device and the TACACS+ server.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.6.3.2 Server Summary

Use the TACACS+ Server Summary page to view and configure information about the TACACS+ Server(s).

To access this page, click **Security > TACACS+ > Server Summary**.




Figure 4.369 Security > TACACS+ > Server Summary

The following table describes the items in the previous figure.

Item	Description
Server	Specifies the TACACS+ Server IP address or Hostname.
Priority	Specifies the order in which the TACACS+ servers are used.
Port	Specifies the authentication port.
Connection Timeout	The amount of time that passes before the connection between the device and the TACACS+ server time out.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new TACACS+ server. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

To add a new TACACS+ server:

Click **Security > TACACS+ > Server Summary > Add**.




Figure 4.370 Security > TACACS+ > Server Summary > Add

The following table describes the items in the previous figure.

Item	Description
Server	Specifies the TACACS+ Server IP address or Hostname.
Priority	Specifies the order in which the TACACS+ servers are used.
Port	Specifies the authentication port.
Key String	Specifies the authentication and encryption key for TACACS+ communications between the device and the TACACS+ server. The key must match the encryption used on the TACACS+ server.
Connection Timeout	The amount of time that passes before the connection between the device and the TACACS+ server time out.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.6.3.3 Server Configuration

Use the TACACS+ Server Configuration page to view and configure information about the TACACS+ Server(s).

To access this page, click **Security > TACACS+ > Server Configuration**.

Server	TestTACACS	
Priority	10	(0 to 65535)
Port	49	(0 to 65535)
Key String	
Connection Timeout	5	(1 to 30 secs)

Submit Remove Refresh Cancel

Figure 4.371 Security > TACACS+ > Server Configuration

The following table describes the items in the previous figure.

Item	Description
Server	Specifies the TACACS+ Server IP address or Hostname.
Priority	Specifies the order in which the TACACS+ servers are used.
Port	Specifies the authentication port.
Key String	Specifies the authentication and encryption key for TACACS+ communications between the device and the TACACS+ server. The key must match the encryption used on the TACACS+ server.
Connection Timeout	The amount of time that passes before the connection between the device and the TACACS+ server time out.
Submit	Click Submit to save the values and update the screen.
Remove	Click Remove to remove the selected entries.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.6.3.4 Source Interface Configuration

Use the TACACS+ Source Interface Configuration page to specify the physical or logical interface to use as the TACACS+ client source interface. When an IP address is configured on the source interface, this address is used for all TACACS+ communications between the local TACACS+ client and the remote TACACS+ server. The IP address of the designated source interface is used in the IP header of TACACS+ management protocol packets. This allows security devices, such as firewalls, to identify all source packets coming from a specific device.

To access this page, click **Security > TACACS+ > Source Interface Configuration**.

Type	<input checked="" type="radio"/> None <input type="radio"/> Interface <input type="radio"/> VLAN <input type="radio"/> Loopback <input type="radio"/> Network <input type="radio"/> Service Port
Interface	Unconfigured
VLAN ID	Unconfigured
Loopback Interface	Unconfigured

Submit Refresh Cancel

Figure 4.372 Security > TACACS+ > Source Interface Configuration

The following table describes the items in the previous figure.

Item	Description
Type	The type of interface to use as the source interface: <ul style="list-style-type: none"> ■ None: The primary IP address of the originating (outbound) interface is used as the source address. ■ Interface: The primary IP address of a physical port is used as the source address. ■ VLAN: The primary IP address of a VLAN routing interface is used as the source address. ■ Loopback: The primary IP address of the loopback interface is used as the source address. A loopback is always reachable, as long as any routing interface is up. ■ Network: The network source IP is used as the source address. ■ Service Port: The management port source IP is used as the source address.
Interface	When the selected Type is Interface, select the physical port to use as the source interface.
VLAN ID	When the selected Type is VLAN, select the VLAN to use as the source interface. The menu contains only the VLAN IDs for VLAN routing interfaces.
Loopback Interface	When the selected Type is Loopback, select the loopback interface to use as the source interface.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.6.4 Authentication Manager

4.6.4.1 Configuration

Use the Authentication Manager Configuration page to control the administrative mode of the Authentication Manager feature, which enables configuration of the sequence and priority of the authentication methods per interface.

To access this page, click **Security > Authentication Manager > Configuration**.




Figure 4.373 Security > Authentication Manager > Configuration

The following table describes the items in the previous figure.

Item	Description
Admin Mode	The administrative mode of the Authentication Manager feature. When Authentication Manager is enabled globally, the authentication methods configured on an interface are executed in the order configured as the device attempts to authenticate clients on that interface.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.6.4.2 Authentication Tiering

Use the Authentication Tiering page to configure the sequence and priority of the authentication methods for the interfaces on the device. When Authentication Manager is enabled globally, the authentication methods configured on an interface are executed in the order configured as the device attempts to authenticate clients on that interface. The default method order is Dot1x, MAC Authentication Bypass (MAB), Captive Portal. To configure the settings for one or more interfaces, select each entry to modify and click Edit. The settings are applied to all selected interfaces. To access this page, click **Security > Authentication Manager > Authentication Tiering**.

Interface	Configured Order	Enabled Order	Configured Priority	Enabled Priority	Authenticated Clients	Re-Authentication Timer
0/1					0	300
0/2					0	300
0/3					0	300
0/4					0	300
0/5					0	300
0/6					0	300
0/7					0	300
0/8					0	300
0/9					0	300
0/10					0	300

Figure 4.374 Security > Authentication Manager > Authentication Tiering

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. When editing information for one or more interfaces, this field identifies the interfaces that are being configured.
Configured Order	The order in which the authentication methods are used to authenticate a client connected to an interface, which can be one or more of the following: <ul style="list-style-type: none"> Dot1x: The port-based authentication method. MAB: MAC Authentication Bypass method that uses the MAC address of the client to determine the kind of network access to provide. Captive Portal: The authentication method that prevents clients from accessing the network until user verification has been established. Captive portal must always be the last method in the list.
Enabled Order	The methods from the list of authentication methods configured on an interface which are administratively enabled in the device.
Configured Priority	The priority of the authentication methods. The default priority of a method is equivalent to its position in the order of the authentication list configured per interface. If the priority of the methods is changed, all clients authenticated using a lower priority method are forced to re-authenticate.
Authenticated Clients	The methods from the list of authentication method priorities configured on an interface which are administratively enabled in the device.
Re-Authentication Timer	Number of clients authenticated on an interface.
Enabled Priority	Interval, in seconds, after which an attempt is made to authenticate an unauthorized port.
Submit	Click Submit to save the values and update the screen.

Item	Description
Edit	Click Edit to edit the selected entries.

4.6.4.3 Authenticated Clients

Use the Authenticated Clients page to view information about the clients connected on the interfaces. If there are no clients connected, the table is empty.

To access this page, click **Security > Authentication Manager > Authenticated Clients**.




Figure 4.375 Security > Authentication Manager > Authenticated Clients

The following table describes the items in the previous figure.

Item	Description
Interface	The local interface associated with the rest of the data in the row.
Logical Interface	The logical port number associated with the client that is connected to the port.
Client MAC Address	The MAC address of the client that is connected to the port.
Authenticated Method	The authentication method used to authenticate a client connected to an interface, which can be one of the following: <ul style="list-style-type: none"> ■ Dot1x: The port-based authentication method. ■ MAB: MAC Authentication Bypass method that uses the MAC address of the client to determine the kind of network access to provide. ■ Captive Portal: The authentication method that prevents clients from accessing the network until user verification has been established.
Authentication State	The current client authentication state, which can be one of the following: <ul style="list-style-type: none"> ■ Success: Indicates authentication succeeded. ■ Failure: Indicated authentication failed.
Authentication Status	The client authentication status, which can be one of the following: <ul style="list-style-type: none"> ■ Authorized: Indicates client is authorized on the port. ■ Unauthorized: Indicates client is not authorized on the port.
Submit	Click Submit to save the values and update the screen.

4.6.4.4 Statistics

Use the Authentication Statistics page to view information about the Authentication Manager client authentication attempts and failures per interface.

To access this page, click **Security > Authentication Manager > Statistics**.

Interface	Dot1x Attempts	Dot1x Failures	MAB Attempts	MAB Failures	Captive Portal Attempts	Captive Portal Failures
0/1	0	0	0	0	0	0
0/2	0	0	0	0	0	0
0/3	0	0	0	0	0	0
0/4	0	0	0	0	0	0
0/5	0	0	0	0	0	0
0/6	0	0	0	0	0	0
0/7	0	0	0	0	0	0
0/8	0	0	0	0	0	0
0/9	0	0	0	0	0	0
0/10	0	0	0	0	0	0

Figure 4.376 Security > Authentication Manager > Statistics

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row.
Dot1x Attempts	The number of attempts made to authenticate a client using the Dot1x authentication method.
Dot1x Failures	The number of attempts that failed when Dot1x method is used for client authentication.
MAB Attempts	The number of attempts made to authenticate a client using the MAC Authentication Bypass (MAB) authentication method.
MAB Failures	The number of attempts that failed when MAB method is used for client authentication.
Captive Portal Failures	The number of attempts made to authenticate a client using the Captive Portal authentication method.
Captive Portal Attempts	The number of attempts that failed when Captive Portal method is used for client authentication.
Submit	Click Submit to save the values and update the screen.
Clear	Click Clear to reset all statistics counters to 0 for the selected interfaces.

4.6.4.5 History

Use the Authentication History page to view the Authentication Manager history log per interface.

To access this page, click **Security > Authentication Manager > History**.

Time Stamp	MAC Address	Authentication Status	Authenticated Method
Table is Empty			

Figure 4.377 Security > Authentication Manager > History

The following table describes the items in the previous figure.

Item	Description
Interface	The menu contains all interfaces in the device. To view the history log on a specific interface, select the interface from the menu.

Item	Description
Time Stamp	The absolute time when the authentication event took place.
MAC Address	The MAC address of the client that is connected to the port.
Authentication Status	The client authentication status, which can be one of the following: <ul style="list-style-type: none"> ■ Authorized: Indicates client is authorized on the port. ■ Unauthorized: Indicates client is not authorized on the port.
Authenticated Method	The authentication method used to authenticate a client connected to an interface, which can be one of the following: <ul style="list-style-type: none"> ■ Dot1x: The port-based authentication method. ■ MAB: MAC Authentication Bypass method that uses the MAC address of the client to determine the kind of network access to provide. ■ Captive Portal: The authentication method that prevents clients from accessing the network until user verification has been established.
Submit	Click Submit to save the values and update the screen.
Clear	Click Clear to clear the Authentication Manager history log on the selected interface.

4.7 QoS

4.7.1 Access Control Lists

4.7.1.1 Summary

Use the Access Control List Summary page to add and remove Access Control Lists (ACLs). ACLs are used to provide traffic flow control, restrict contents of routing updates, decide which types of traffic are forwarded or blocked, and above all provide security for the network. There are three main steps to configuring an ACL:

1. Create an ACL. (Use the current page.)
2. Add rules to the ACL and configure the rule criteria. (Use the Access Control List Configuration page.)
3. Apply the ACL to one or more interfaces. (Use the Access Control List Interface Summary page.)

To access this page, click **QoS > Access Control Lists > Summary**.




Figure 4.378 QoS > Access Control Lists > Summary

The following table describes the items in the previous figure.

Item	Description
ACL Identifier	The name or number that identifies the ACL. The permitted identifier depends on the ACL type. Standard and Extended IPv4 ACLs use numbers within a set range, and Named IPv4, IPv6, and MAC ACLs use alphanumeric characters. The ID of a Named IPv4 ACL must begin with a letter, and not a number.

Item	Description
ACL Type	<p>The type of ACL. The ACL type determines the criteria that can be used to match packets. The type also determines which attributes can be applied to matching traffic. IPv4 ACLs classify Layer 3 and Layer 4 IPv4 traffic, IPv6 ACLs classify Layer 3 and Layer 4 IPv6 traffic, and MAC ACLs classify Layer 2 traffic. The ACL types are as follows:</p> <ul style="list-style-type: none"> ■ IPv4 Standard: Match criteria is based on the source address of IPv4 packets. ■ IPv4 Extended: Match criteria can be based on the source and destination addresses, source and destination Layer 4 ports, and protocol type of IPv4 packets. ■ IPv4 Named: Match criteria is the same as IPv4 Extended ACLs, but the ACL ID can be an alphanumeric name instead of a number. ■ IPv6 Named: Match criteria can be based on information including the source and destination IPv6 addresses, source and destination Layer 4 ports, and protocol type within IPv6 packets. ■ Extended MAC: Match criteria can be based on the source and destination MAC addresses, 802.1p user priority, VLAN ID, and EtherType value within Ethernet frames.
Rules Used	The number of rules currently configured for the ACL.
Direction	Indicates whether the packet is checked against the rules in an ACL when it is received on an interface (Inbound) or after it has been received, routed, and is ready to exit an interface (Outbound).
Interface	Each interface to which the ACL has been applied.
VLAN	Each VLAN to which the ACL has been applied.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new ACL.
Edit	Click Edit to edit the selected entries. See the following procedure.
Remove	Click Remove to remove the selected entries.

To add a new ACL:

Click **QoS > Access Control Lists > Summary > Add**.




Figure 4.379 QoS > Access Control Lists > Summary > Add

The following table describes the items in the previous figure.

Item	Description
ACL Type	<p>The type of ACL. The ACL type determines the criteria that can be used to match packets. The type also determines which attributes can be applied to matching traffic. IPv4 ACLs classify Layer 3 and Layer 4 IPv4 traffic, IPv6 ACLs classify Layer 3 and Layer 4 IPv6 traffic, and MAC ACLs classify Layer 2 traffic. The ACL types are as follows:</p> <ul style="list-style-type: none"> ■ IPv4 Standard: Match criteria is based on the source address of IPv4 packets. ■ IPv4 Extended: Match criteria can be based on the source and destination addresses, source and destination Layer 4 ports, and protocol type of IPv4 packets. ■ IPv4 Named: Match criteria is the same as IPv4 Extended ACLs, but the ACL ID can be an alphanumeric name instead of a number. ■ IPv6 Named: Match criteria can be based on information including the source and destination IPv6 addresses, source and destination Layer 4 ports, and protocol type within IPv6 packets. ■ Extended MAC: Match criteria can be based on the source and destination MAC addresses, 802.1p user priority, VLAN ID, and EtherType value within Ethernet frames.
ACL Identifier	<p>The name or number that identifies the ACL. The permitted identifier depends on the ACL type. Standard and Extended IPv4 ACLs use numbers within a set range, and Named IPv4, IPv6, and MAC ACLs use alphanumeric characters. The ID of a Named IPv4 ACL must begin with a letter, and not a number.</p>
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.7.1.2 Configuration

Use the Access Control List Configuration page to configure rules for the existing Access Control Lists (ACLs) on the system and to view summary information about the rules that have been added to an ACL. Each ACL rule is configured to match one or more aspects of traffic on the network. When a packet matches the conditions in a rule, it is handled according to the configured action (permit or deny) and attributes. Each ACL can have multiple rules, but the final rule for every ACL is an implicit deny all rule. For each rule, a packet must match all the specified criteria in order for the specified rule action (Permit/Deny) to take place.

To access this page, click **QoS > Access Control Lists > Configuration**.

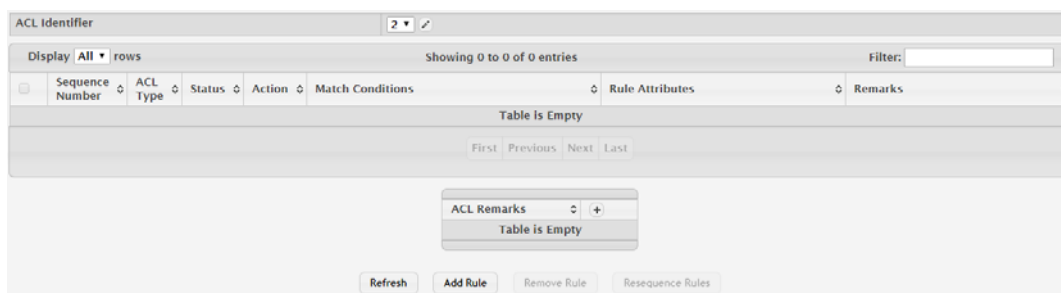


Figure 4.380 QoS > Access Control Lists > Configuration

The following table describes the items in the previous figure.

Item	Description
ACL Identifier	The menu contains the ID for each ACL that exists on the system. Before you add or remove a rule, you must select the ID of the ACL from the menu. For ACLs with alphanumeric names, click the Edit icon to change the ACL ID. The ID of a named ACL must begin with a letter, and not a number. The ACL identifier for IPv4 Standard and IPv4 Extended ACLs cannot be changed.
Sequence Number	The number that indicates the position of a rule within the ACL. If the sequence number is not specified during rule creation, the rule is automatically assigned a sequence number after it is successfully added to the ACL. The rules are displayed based on their position within the ACL, but can also be renumbered. Packets are checked against the rule criteria in order, from the lowest-numbered rule to the highest. When the packet matches the criteria in a rule, it is handled according to the rule action and attributes. If no rule matches a packet, the packet is discarded based on the implicit deny all rule, which is the final rule in every ACL.
ACL Type	<p>The type of ACL. The ACL type determines the criteria that can be used to match packets. The type also determines which attributes can be applied to matching traffic. IPv4 ACLs classify Layer 3 and Layer 4 IPv4 traffic, IPv6 ACLs classify Layer 3 and Layer 4 IPv6 traffic, and MAC ACLs classify Layer 2 traffic. The ACL types are as follows:</p> <ul style="list-style-type: none"> ■ IPv4 Standard: Match criteria is based on the source address of IPv4 packets. ■ IPv4 Extended: Match criteria can be based on the source and destination addresses, source and destination Layer 4 ports, and protocol type of IPv4 packets. ■ IPv4 Named: Match criteria is the same as IPv4 Extended ACLs, but the ACL ID can be an alphanumeric name instead of a number. ■ IPv6 Named: Match criteria can be based on information including the source and destination IPv6 addresses, source and destination Layer 4 ports, and protocol type within IPv6 packets. ■ Extended MAC: Match criteria can be based on the source and destination MAC addresses, 802.1p user priority, VLAN ID, and EtherType value within Ethernet frames.
Status	Indicates whether the ACL is active. If the ACL is a time-based ACL that includes a time range, the ACL is active only during the periods specified within the time range. If an ACL does not include a time range, the status is always active.
Action	<p>The action to take when a packet or frame matches the criteria in the rule:</p> <ul style="list-style-type: none"> ■ Permit: The packet or frame is forwarded. ■ Deny: The packet or frame is dropped. <p><i>NOTE: When configuring ACL rules in the Add Access Control List Rule window, the selected action determines which fields can be configured. Not all fields are available for both Permit and Deny actions.</i></p>
Match Conditions	The criteria used to determine whether a packet or frame matches the ACL rule.
Rule Attributes	Each action - beyond the basic Permit and Deny actions - to perform on the traffic that matches the rule.

Item	Description
Remarks	One or more remarks configured for the selected ACL and associated with the rule during rule creation. To delete a remark associated with the rule, click the: (minus) button preceding remark. You must confirm the action before the rule associated remark is removed.
Refresh	Click Refresh to update the screen.
Add Rule	Click Add Rule to add a new ACL rule. See the following procedure.
Remove Rule	Click Remove Rule to remove the selected entries.
Resequence Rules	Click Resequence Rules to resequence rules for the selected entries. See the following procedure.

To add a new ACL rule:

Click **QoS > Access Control Lists > Configuration > Add Rule**.

Figure 4.381 QoS > Access Control Lists > Configuration > Add Rule

The following table describes the items in the previous figure.

Item	Description
Sequence Number	The number that indicates the position of a rule within the ACL. If the sequence number is not specified during rule creation, the rule is automatically assigned a sequence number after it is successfully added to the ACL. The rules are displayed based on their position within the ACL, but can also be renumbered. Packets are checked against the rule criteria in order, from the lowest-numbered rule to the highest. When the packet matches the criteria in a rule, it is handled according to the rule action and attributes. If no rule matches a packet, the packet is discarded based on the implicit deny all rule, which is the final rule in every ACL.
Action	The action to take when a packet or frame matches the criteria in the rule: <ul style="list-style-type: none"> ■ Permit: The packet or frame is forwarded. ■ Deny: The packet or frame is dropped. <p><i>NOTE: When configuring ACL rules in the Add Access Control List Rule window, the selected action determines which fields can be configured. Not all fields are available for both Permit and Deny actions.</i></p>
Match Criteria (IPv4 ACLs)	
Every	When this option is selected, all packets will match the rule and will be either permitted or denied. This option is exclusive to all other match criteria, so if Every is selected, no other match criteria can be configured. To configure specific match criteria, this option must be clear.

Item	Description
Protocol	The IANA-assigned protocol number to match within the IP packet. You can also specify one of the following keywords: EIGRP, GRE, ICMP, IGMP, IP, IPIP, OSPF, PIM, TCP, or UDP. The function is only available for IPv4 Extended and IPv4 Named ACLs.
Fragments	IP ACL rule to match on fragmented IP packets. The function is only available for IPv4 Extended and IPv4 Named ACLs.
Source IP Address / Wildcard Mask	The source port IP address in the packet and source IP wildcard mask (in the second field) to compare to the IP address in a packet header. Wild card masks determines which bits in the IP address are used and which bits are ignored. A wild card mask of 255.255.255.255 indicates that no bit is important. A wildcard of 0.0.0.0 indicates that all of the bits are important. For example, enter a wildcard mask of 0.0.0.0 to specify a host. Wildcard masking for ACLs operates differently from a subnet mask. A wildcard mask is in essence the inverse of a subnet mask. With a subnet mask, the mask has ones (1's) in the bit positions that are used for the network address, and has zeros (0's) for the bit positions that are not used. In contrast, a wildcard mask has (0's) in a bit position that must be checked. A '1' in a bit position of the ACL mask indicates the corresponding bit can be ignored. This field is required when you configure a source IP address.
Source L4 Port	The TCP/UDP source port to match in the packet header. Select one of the following options: Equal, Not Equal, Less Than, Greater Than, or Range and specify the port number or keyword. TCP port keywords include BGP, Domain, Echo, FTP, FTP Data, HTTP, SMTP, Telnet, WWW, POP2, and POP3. UDP port keywords include Domain, Echo, NTP, RIP, SNMP, TFTP, TIME, and WHO. The function is only available for IPv4 Extended and IPv4 Named ACLs.
Destination IP Address / Wildcard Mask	The destination port IP address in the packet and destination IP wildcard mask (in the second field) to compare to the IP address in a packet header. Wild card masks determines which bits in the IP address are used and which bits are ignored. A wild card mask of 255.255.255.255 indicates that no bit is important. A wildcard of 0.0.0.0 indicates that all of the bits are important. For example, enter a wildcard mask of 0.0.0.0 to specify a host. Wildcard masking for ACLs operates differently from a subnet mask. A wildcard mask is in essence the inverse of a subnet mask. With a subnet mask, the mask has ones (1's) in the bit positions that are used for the network address, and has zeros (0's) for the bit positions that are not used. In contrast, a wildcard mask has (0's) in a bit position that must be checked. A 1 in a bit position of the ACL mask indicates the corresponding bit can be ignored. This field is required when you configure a destination IP address.
Destination L4 Port	The TCP/UDP destination port to match in the packet header. Select one of the following options: Equal, Not Equal, Less Than, Greater Than, or Range and specify the port number or keyword. TCP port keywords include BGP, Domain, Echo, FTP, FTP Data, HTTP, SMTP, Telnet, WWW, POP2, and POP3. UDP port keywords include Domain, Echo, NTP, RIP, SNMP, TFTP, TIME, and WHO. The function is only available for IPv4 Extended and IPv4 Named ACLs.
TTL Field Value	IP ACL rule to match on the specified TTL field value. The function is only available for IPv4 Extended and IPv4 Named ACLs.
IGMP Type	IP ACL rule to match on the specified IGMP message type. This option is available only if the protocol is IGMP. The function is only available for IPv4 Extended and IPv4 Named ACLs.

Item	Description
ICMP Type	IP ACL rule to match on the specified ICMP message type. This option is available only if the protocol is ICMP. The function is only available for IPv4 Extended and IPv4 Named ACLs.
ICMP Code	IP ACL rule to match on the specified ICMP message code. This option is available only if the protocol is ICMP. The function is only available for IPv4 Extended and IPv4 Named ACLs.
ICMP Message	IP ACL rule to match on the ICMP message type and code. Specify one of the following supported ICMP messages: Echo, Echo-Reply, Host-Redirect, Mobile-Redirect, Net-Redirect, Net-Unreachable, Redirect, Packet-Too-Big, Port-Unreachable, Source-Quench, Router-Solicitation, Router-Advertisement, Time-Exceeded, TTL-Exceeded, and Unreachable. This option is available only if the protocol is ICMP. The function is only available for IPv4 Extended and IPv4 Named ACLs.
TCP Flags	IP ACL rule to match on the TCP flags. When a + flag is specified, a match occurs if the flag is set in the TCP header. When a - flag is specified, a match occurs if the flag is not set in the TCP header. When Established is specified, a match occurs if either RST or ACK bits are set in the TCP header. This option is available only if the protocol is TCP. The function is only available for IPv4 Extended and IPv4 Named ACLs.
Service Type	<p>The service type to match in the IP header. The options in this menu are alternative ways of specifying a match condition for the same Service Type field in the IP header, but each service type uses a different user notation. After you select the service type, specify the value for the service type in the appropriate field. Only the field associated with the selected service type can be configured. The function is only available for IPv4 Extended and IPv4 Named ACLs. The services types are as follows:</p> <ul style="list-style-type: none"> ■ IP DSCP: Matches the packet IP DiffServ Code Point (DSCP) value to the rule. The DSCP value is defined as the high-order six bits of the Service Type octet in the IP header. ■ IP Precedence: Matches the IP Precedence value to the rule. The IP Precedence field in a packet is defined as the high-order three bits of the Service Type octet in the IP header. ■ IP TOS Bits: Matches on the Type of Service (TOS) bits in the IP header. The IP TOS field in a packet is defined as all eight bits of the Service Type octet in the IP header. For example, to check for an IP TOS value having bits 7 and 5 set and bit 1 clear, where bit 7 is most significant, use a TOS Bits value of 0xA0 and a TOS Mask of 0xFF. <ul style="list-style-type: none"> – TOS Bits: Requires the bits in a packet's TOS field to match the two-digit hexadecimal number entered in this field. – TOS Mask: The bit positions that are used for comparison against the IP TOS field in a packet. Specifying TOS Mask is optional.
Match Criteria (IPv6 ACLs)	
Every	When this option is selected, all packets will match the rule and will be either permitted or denied. This option is exclusive to all other match criteria, so if Every is selected, no other match criteria can be configured. To configure specific match criteria, this option must be clear.
Protocol	The IANA-assigned protocol number to match within the IP packet. You can also specify one of the following keywords: ICMPv6, IPv6, TCP, or UDP.
Fragments	IPv6 ACL rule to match on fragmented IP packets.

Item	Description
Source Prefix / Prefix Length	The IPv6 prefix combined with IPv6 prefix length of the network or host from which the packet is being sent. To indicate a destination host, specify an IPv6 prefix length of 128.
Source L4 Port	The TCP/UDP source port to match in the packet header. Select one of the following options: Equal, Not Equal, Less Than, Greater Than, or Range and specify the port number or keyword. TCP port keywords include BGP, Domain, Echo, FTP, FTP Data, HTTP, SMTP, Telnet, WWW, POP2, and POP3. UDP port keywords include Domain, Echo, NTP, RIP, SNMP, TFTP, TIME, and WHO.
Destination Prefix / Prefix Length	The IPv6 prefix combined with the IPv6 prefix length to be compared to a packet's destination IPv6 address as a match criteria for the IPv6 ACL rule. To indicate a destination host, specify an IPv6 prefix length of 128.
Destination L4 Port	The TCP/UDP destination port to match in the packet header. Select one of the following options: Equal, Not Equal, Less Than, Greater Than, or Range and specify the port number or keyword. TCP port keywords include BGP, Domain, Echo, FTP, FTP Data, HTTP, SMTP, Telnet, WWW, POP2, and POP3. UDP port keywords include Domain, Echo, NTP, RIP, SNMP, TFTP, TIME, and WHO.
TTL Field Value	IP ACL rule to match on the specified TTL field value. The function is only available for IPv6 Named ACLs.
ICMP Type	IPv6 ACL rule to match on the specified ICMP message type. This option is available only if the protocol is ICMPv6.
ICMP Code	IPv6 ACL rule to match on the specified ICMP message code. This option is available only if the protocol is ICMPv6.
ICMP Message	IPv6 ACL rule to match on the ICMP message type and code. Specify one of the following supported ICMPv6 messages: Destination-Unreachable, Echo-Request, Echo-Reply, Header, Hop-Limit, MLD-Query, MLD-Reduction, MLD-Report, ND-NA, ND-NS, Next-Header, No-Admin, No-Route, Packet-Too-Big, Port-Unreachable, Router-Solicitation, Router-Advertisement, Router-Renumbering, Time-Exceeded, and Unreachable. This option is available only if the protocol is ICMPv6.
TCP Flags	IPv6 ACL rule to match on the TCP flags. When a + flag is specified, a match occurs if the flag is set in the TCP header. When a - flag is specified, a match occurs if the flag is not set in the TCP header. When Established is specified, a match occurs if either RST or ACK bits are set in the TCP header. This option is available only if the protocol is TCP.
Flow Label	A 20-bit number that is unique to an IPv6 packet, used by end stations to signify quality-of-service handling in routers.
IP DSCP	The IP DSCP value in the IPv6 packet to match to the rule. The DSCP value is defined as the high-order six bits of the Service Type octet in the IPv6 header.
Routing	IPv6 ACL rule to match on routed packets.
Match Criteria (MAC ACLs)	
Every	When this option is selected, all packets will match the rule and will be either permitted or denied. This option is exclusive to all other match criteria, so if Every is selected, no other match criteria can be configured. To configure specific match criteria, this option must be clear.
CoS	The 802.1p user priority value to match within the Ethernet frame.

Item	Description
Ethertype	The EtherType value to match in an Ethernet frame. Specify the number associated with the EtherType or specify one of the following keywords: AppleTalk, ARP, IBM SNA, IPv4, IPv6, IPX, MPLS, Unicast, NETBIOS, NOVELL, PPPoE, or RARP.
Source MAC Address / Mask	The MAC address to match to an Ethernet frame's source port MAC address. If desired, enter the MAC Mask associated with the source MAC to match. The MAC address mask specifies which bits in the source MAC to compare against an Ethernet frame. Use F's and zeros in the MAC mask, which is in a wildcard format. An F means that the bit is not checked, and a zero in a bit position means that the data must equal the value given for that bit. For example, if the MAC address is aa:bb:cc:dd:ee:ff, and the mask is 00:00:ff:ff:ff:ff, all MAC addresses with aa:bb:xx:xx:xx:xx result in a match (where x is any hexadecimal number).
Destination MAC Address / Mask	The MAC address to match to an Ethernet frame's destination port MAC address. If desired, enter the MAC Mask associated with the destination MAC to match. The MAC address mask specifies which bits in the destination MAC to compare against an Ethernet frame. Use F's and zeros in the MAC mask, which is in a wildcard format. An F means that the bit is not checked, and a zero in a bit position means that the data must equal the value given for that bit. For example, if the MAC address is aa:bb:cc:dd:ee:ff, and the mask is 00:00:ff:ff:ff:ff, all MAC addresses with aa:bb:xx:xx:xx:xx result in a match (where x is any hexadecimal number).
VLAN	The VLAN ID to match within the Ethernet frame.
Rule Attributes	
Assign Queue	The number that identifies the hardware egress queue that will handle all packets matching this rule.
Interface	The interface to use for the action: <ul style="list-style-type: none"> ■ Redirect: Allows traffic that matches a rule to be redirected to the selected interface instead of being processed on the original port. The redirect function and mirror function are mutually exclusive. ■ Mirror: Provides the ability to mirror traffic that matches a rule to the selected interface. Mirroring is similar to the redirect function, except that in flow-based mirroring a copy of the permitted traffic is delivered to the mirror interface while the packet itself is forwarded normally through the device.
Log	When this option is selected, logging is enabled for this ACL rule (subject to resource availability in the device). If the Access List Trap Flag is also enabled, this will cause periodic traps to be generated indicating the number of times this rule went into effect during the current report interval. A fixed 5 minute report interval is used for the entire system. A trap is not issued if the ACL rule hit count is zero for the current interval.
Time Range Name	The name of the time range that will impose a time limitation on the ACL rule. If a time range with the specified name does not exist, and the ACL containing this ACL rule is associated with an interface, the ACL rule is applied immediately. If a time range with specified name exists, and the ACL containing this ACL rule is associated with an interface, the ACL rule is applied when the time-range with specified name becomes active. The ACL rule is removed when the time-range with specified name becomes inactive.
Committed Rate / Burst Size	The allowed transmission rate for packets on the interface (Committed Rate), and the number of bytes allowed in a temporary traffic burst (Burst Rate).

Item	Description
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

To resequence rules for an ACL:

Click **QoS > Access Control Lists > Configuration > Add Rule**.

Figure 4.382 QoS > Access Control Lists > Configuration > Add Rule

The following table describes the items in the previous figure.

Item	Description
Sequence Start	The starting sequence number for resequencing the existing rules.
Sequence Step	The increment of sequence numbers for resequencing the existing rules.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.7.1.3 Interfaces

Use the Access Control List Interface Summary page to associate one or more ACLs with one or more interfaces on the device. When an ACL is associated with an interface, traffic on the port is checked against the rules defined within the ACL until a match is found. If the traffic does not match any rules within an ACL, it is dropped because of the implicit deny all rule at the end of each ACL.

To access this page, click **QoS > Access Control Lists > Interfaces**.

Figure 4.383 QoS > Access Control Lists > Interfaces

The following table describes the items in the previous figure.

Item	Description
Interface	The interface that has an associated ACL.
Direction	Indicates whether the packet is checked against the rules in an ACL when it is received on an interface (Inbound) or after it has been received, routed, and is ready to exit an interface (Outbound).
Sequence Number	The order the ACL is applied to traffic on the interface relative to other ACLs associated with the interface in the same direction. When multiple ACLs are applied to the same interface in the same direction, the ACL with the lowest sequence number is applied first, and the other ACLs are applied in ascending numerical order.

Item	Description
ACL Type	The type of ACL. The ACL type determines the criteria that can be used to match packets. The type also determines which attributes can be applied to matching traffic. IPv4 ACLs classify Layer 3 and Layer 4 IPv4 traffic, IPv6 ACLs classify Layer 3 and Layer 4 IPv6 traffic, and MAC ACLs classify Layer 2 traffic. The ACL types are as follows: <ul style="list-style-type: none"> ■ IPv4 Standard: Match criteria is based on the source address of IPv4 packets. ■ IPv4 Extended: Match criteria can be based on the source and destination addresses, source and destination Layer 4 ports, and protocol type of IPv4 packets. ■ IPv4 Named: Match criteria is the same as IPv4 Extended ACLs, but the ACL ID can be an alphanumeric name instead of a number. ■ IPv6 Named: Match criteria can be based on information including the source and destination IPv6 addresses, source and destination Layer 4 ports, and protocol type within IPv6 packets. ■ Extended MAC: Match criteria can be based on the source and destination MAC addresses, 802.1p user priority, VLAN ID, and EtherType value within Ethernet frames.
ACL Identifier	The name or number that identifies the ACL. When applying an ACL to an interface, the ACL Identifier menu includes only the ACLs within the selected ACL Type.
Refresh	Click Refresh to update the screen.
Add	Click Add to apply an ACL to an interface. See the following procedure.
Remove	Click Remove to remove the association between an interface and an ACL.

To apply an ACL to an interface:

Click **QoS > Access Control Lists > Interfaces > Add**.

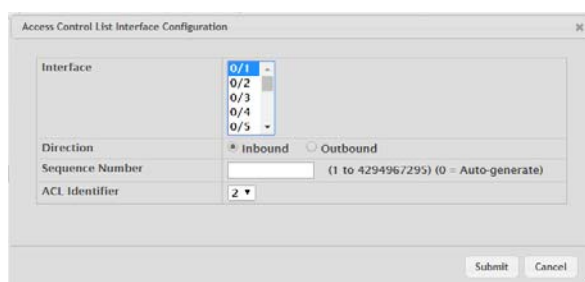


Figure 4.384 QoS > Access Control Lists > Interfaces > Add

The following table describes the items in the previous figure.

Item	Description
Interface	The interface that has an associated ACL.
Direction	Indicates whether the packet is checked against the rules in an ACL when it is received on an interface (Inbound) or after it has been received, routed, and is ready to exit an interface (Outbound).
Sequence Number	The order the ACL is applied to traffic on the interface relative to other ACLs associated with the interface in the same direction. When multiple ACLs are applied to the same interface in the same direction, the ACL with the lowest sequence number is applied first, and the other ACLs are applied in ascending numerical order.

Item	Description
ACL Identifier	The name or number that identifies the ACL. When applying an ACL to an interface, the ACL Identifier menu includes only the ACLs within the selected ACL Type.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.7.1.4 VLANs

Use the Access Control List VLAN Summary page to associate one or more ACLs with one or more VLANs on the device.

To access this page, click **QoS > Access Control Lists > VLANs**.




Figure 4.385 QoS > Access Control Lists > VLANs

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The ID of the VLAN associated with the rest of the data in the row. When associating a VLAN with an ACL, use this field to select the desired VLAN.
Direction	Indicates whether the packet is checked against the rules in an ACL when it is received on a VLAN (Inbound) or after it has been received, routed, and is ready to exit a VLAN (Outbound).
Sequence Number	The order the ACL is applied to traffic on the VLAN relative to other ACLs associated with the VLAN in the same direction. When multiple ACLs are applied to the same VLAN in the same direction, the ACL with the lowest sequence number is applied first, and the other ACLs are applied in ascending numerical order.
ACL Type	The type of ACL. The ACL type determines the criteria that can be used to match packets. The type also determines which attributes can be applied to matching traffic. IPv4 ACLs classify Layer 3 and Layer 4 IPv4 traffic, IPv6 ACLs classify Layer 3 and Layer 4 IPv6 traffic, and MAC ACLs classify Layer 2 traffic. The ACL types are as follows: <ul style="list-style-type: none"> ■ IPv4 Standard: Match criteria is based on the source address of IPv4 packets. ■ IPv4 Extended: Match criteria can be based on the source and destination addresses, source and destination Layer 4 ports, and protocol type of IPv4 packets. ■ IPv4 Named: Match criteria is the same as IPv4 Extended ACLs, but the ACL ID can be an alphanumeric name instead of a number. ■ IPv6 Named: Match criteria can be based on information including the source and destination IPv6 addresses, source and destination Layer 4 ports, and protocol type within IPv6 packets. ■ Extended MAC: Match criteria can be based on the source and destination MAC addresses, 802.1p user priority, VLAN ID, and EtherType value within Ethernet frames.

Item	Description
ACL Identifier	The name or number that identifies the ACL. The permitted identifier depends on the ACL type. Standard and Extended IPv4 ACLs use numbers within a set range, and Named IPv4, IPV6, and MAC ACLs use alphanumeric characters.
Refresh	Click Refresh to update the screen.
Add	Click Add to associate an ACL with a VLAN. See the following procedure.
Remove	Click Remove to remove the association between a VLAN and an ACL.

To associate an ACL with a VLAN:

Click **QoS > Access Control Lists > VLANs > Add**.

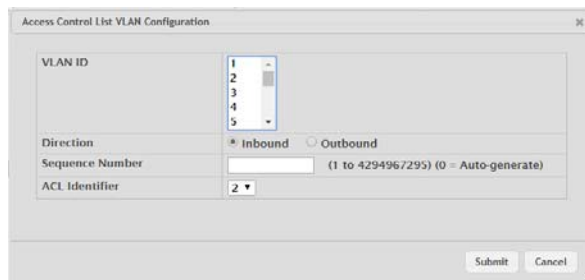


Figure 4.386 QoS > Access Control Lists > VLANs > Add

The following table describes the items in the previous figure.

Item	Description
VLAN ID	The ID of the VLAN associated with the rest of the data in the row. When associating a VLAN with an ACL, use this field to select the desired VLAN.
Direction	Indicates whether the packet is checked against the rules in an ACL when it is received on a VLAN (Inbound) or after it has been received, routed, and is ready to exit a VLAN (Outbound).
Sequence Number	The order the ACL is applied to traffic on the VLAN relative to other ACLs associated with the VLAN in the same direction. When multiple ACLs are applied to the same VLAN in the same direction, the ACL with the lowest sequence number is applied first, and the other ACLs are applied in ascending numerical order.
ACL Identifier	The name or number that identifies the ACL. The permitted identifier depends on the ACL type. Standard and Extended IPv4 ACLs use numbers within a set range, and Named IPv4, IPV6, and MAC ACLs use alphanumeric characters
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.7.1.5 Control Plane

Use the Access Control List Control Plane Configuration page to define controlled management access to the device. Control plane ACLs allow you to determine which addresses or protocols are allowed to access the management interface on the device. The control plane ACLs are applied to management access through the in-band (production network) ports only. Inbound traffic on the CPU port is checked against the rules defined within the ACL until a match is found. If the traffic does not match any rules within an ACL, it is dropped because of the implicit deny all rule at the end of each ACL.

To access this page, click **QoS > Access Control Lists > Control Plane**.




Figure 4.387 QoS > Access Control Lists > Control Plane

The following table describes the items in the previous figure.

Item	Description
ACL Identifier	The name or number that identifies the ACL.
ACL Type	The type of ACL. The ACL type determines the criteria that can be used to match packets. The type also determines which attributes can be applied to matching traffic. IPv4 ACLs classify Layer 3 and Layer 4 IPv4 traffic, IPv6 ACLs classify Layer 3 and Layer 4 IPv6 traffic, and MAC ACLs classify Layer 2 traffic. The ACL types are as follows: <ul style="list-style-type: none"> ■ IPv4 Standard: Match criteria is based on the source address of IPv4 packets. ■ IPv4 Extended: Match criteria can be based on the source and destination addresses, source and destination Layer 4 ports, and protocol type of IPv4 packets. ■ IPv4 Named: Match criteria is the same as IPv4 Extended ACLs, but the ACL ID can be an alphanumeric name instead of a number. ■ IPv6 Named: Match criteria can be based on information including the source and destination IPv6 addresses, source and destination Layer 4 ports, and protocol type within IPv6 packets. ■ Extended MAC: Match criteria can be based on the source and destination MAC addresses, 802.1p user priority, VLAN ID, and EtherType value within Ethernet frames.
Sequence Number	The order the ACL is applied to traffic on the interface relative to other ACLs associated with the interface in the same direction. When multiple ACLs are applied to the same interface in the same direction, the ACL with the lowest sequence number is applied first, and the other ACLs are applied in ascending numerical order.
Refresh	Click Refresh to update the screen.
Add	Click Add to apply an ACL to the CPU interface. See the following procedure.
Remove	Click Remove to remove the association between the CPU interface and an ACL.

To apply an ACL to the CPU interface:

Click **QoS > Access Control Lists > Control Plane > Add**.

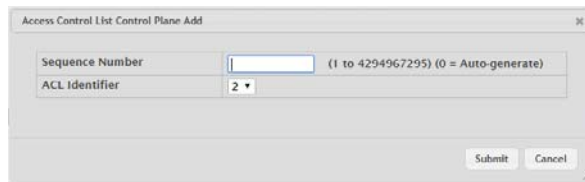


Figure 4.388 QoS > Access Control Lists > Control Plane > Add

The following table describes the items in the previous figure.

Item	Description
Sequence Number	The order the ACL is applied to traffic on the interface relative to other ACLs associated with the interface in the same direction. When multiple ACLs are applied to the same interface in the same direction, the ACL with the lowest sequence number is applied first, and the other ACLs are applied in ascending numerical order.
ACL Identifier	The name or number that identifies the ACL.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.7.1.6 Statistics

The Access Control List Statistics page displays the statistical information about the packets forwarded or discarded by the port that match the configured rules within an Access Control List (ACL). Each ACL rule is configured to match one or more aspects of traffic on the network. When a packet matches the conditions in a rule, the counter associated with the rule gets incremented, until it reaches the rollover value of the counter. ACL counters do not interact with DiffServ policies or Policy-based Routing counters.

To access this page, click **QoS > Access Control Lists > Statistics**.



Figure 4.389 QoS > Access Control Lists > Statistics

The following table describes the items in the previous figure.

Item	Description
ACL Type	<p>The type of ACL. The ACL type determines the criteria that can be used to match packets. The type also determines which attributes can be applied to matching traffic. IPv4 ACLs classify Layer 3 and Layer 4 IPv4 traffic, IPv6 ACLs classify Layer 3 and Layer 4 IPv6 traffic, and MAC ACLs classify Layer 2 traffic. The ACL types are as follows:</p> <ul style="list-style-type: none"> ■ IPv4 Standard: Match criteria is based on the source address of the IPv4 packets. ■ IPv4 Extended: Match criteria can be based on the source and destination addresses, source and destination Layer 4 ports, and protocol type of the IPv4 packets. ■ IPv4 Named: Match criteria is the same as IPv4 Extended ACLs, but the ACL ID can be an alphanumeric name instead of a number. ■ IPv6 Named: Match criteria can be based on information including the source and destination IPv6 addresses, source and destination Layer 4 ports, and protocol type within the IPv6 packets. ■ Extended MAC: Match criteria can be based on the source and destination MAC addresses, 802.1p user priority, VLAN ID, and EtherType value within the Ethernet frames.
ACL Identifier	<p>A list of ACL IDs that exist on the system for a given ACL type. To view the rule(s) within an ACL, you must select the ID of the ACL from the list. The ACL rules are not displayed when option All is selected. Option All lets you clear the hit count for an ACL type.</p>
Sequence Number	<p>The number that indicates the position of a rule within the ACL.</p>
Action	<p>The action to take when a packet or frame matches the criteria in the rule:</p> <ul style="list-style-type: none"> ■ Permit: The packet or frame is forwarded. ■ Deny: The packet or frame is dropped.
Match Conditions	<p>The criteria used to determine whether a packet or frame matches the ACL rule.</p>
Rule Attributes	<p>Each action - beyond the basic Permit and Deny actions - to perform on the traffic that matches the rule.</p>
Hit Count	<p>Indicates the number of packets that match the configured rule in an ACL. If a rule is configured without rate limit, then the hit count is the number of matched packets forwarded or discarded by the port. If a rule is configured with rate limit, then if the sent traffic rate exceeds the configured rate, the hit count displays the matched packet count equal to the sent rate, despite packets getting dropped beyond the configured limit. If the sent traffic rate is less than the configured rate, the hit count displays only the matched packet count.</p>
Refresh	<p>Click Refresh to update the screen.</p>
Clear Rule Counter	<p>Click Clear Rule Counter to clear the hit count for one or more configured rules within an ACL.</p>
Clear ACL Counters	<p>Click Clear ACL Counters to remove the association between an interface and an ACL.</p>

4.7.2 Auto VoIP

4.7.2.1 Global

Use the Auto VoIP Global Configuration page to configure the VLAN ID for the Auto VoIP VLAN or to reset the current Auto VoIP VLAN ID to the default value. Voice over Internet Protocol (VoIP) enables telephone calls over a data network. Because voice traffic is typically more time-sensitive than data traffic, the Auto VoIP feature helps provide a classification mechanism for voice packets so that they can be prioritized above data packets in order to provide better Quality of Service (QoS). With the Auto VoIP feature, voice prioritization is provided based on call-control protocols (SIP, SCCP, H.323) and/or OUI bits. When the device identifies voice traffic, it is placed in the VLAN specified on this page. The Auto VoIP feature does not rely on LLDP-MED support in connected devices.

To access this page, click **QoS > Auto VoIP > Global**.

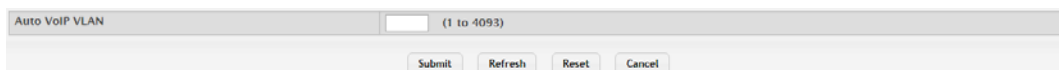


Figure 4.390 QoS > Auto VoIP > Global

The following table describes the items in the previous figure.

Item	Description
Auto VoIP VLAN	The VLAN used to segregate VoIP traffic from other non-voice traffic.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Reset	Click Reset to reset the voice VLAN to the default value.
Cancel	Click Cancel to restore default value.

4.7.2.2 OUI Table

Use the OUI Table Summary page to add and remove Organizationally Unique Identifiers (OUIs) from the OUI database the device maintains. Device hardware manufacturers can include an OUI in a network adapter to help identify the device. The OUI is a unique 24-bit number assigned by the IEEE registration authority. Several default OUIs have been preconfigured in the OUI database on the device.

To access this page, click **QoS > Auto VoIP > OUI Table**.

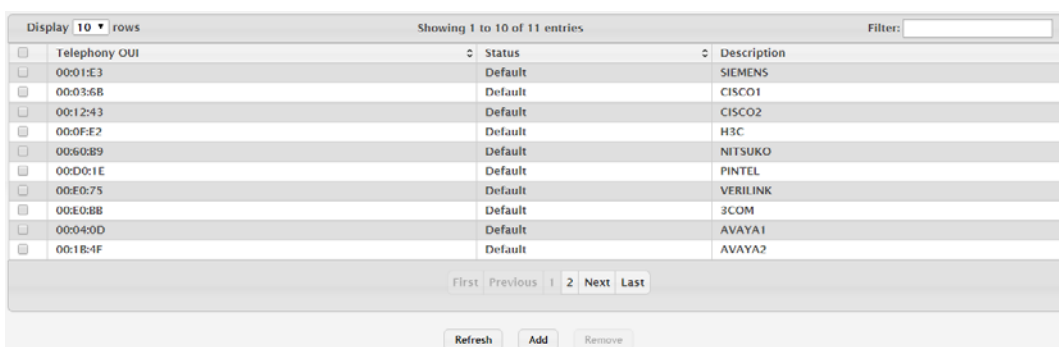


Figure 4.391 QoS > Auto VoIP > OUI Table

The following table describes the items in the previous figure.

Item	Description
Telephony OUI	The unique OUI that identifies the device manufacturer or vendor. The OUI is specified in three octet values (each octet is represented as two hexadecimal digits) separated by colons.

Item	Description
Status	Identifies whether the OUI is preconfigured on the system (Default) or added by a user (Configured).
Description	Identifies the manufacturer or vendor associated with the OUI.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new OUI. See the following procedure.
Remove	Click Remove to remove the selected entries.

To add a new OUI:

Click **QoS > Auto VoIP > OUI Table > Add**.

Figure 4.392 QoS > Auto VoIP > OUI Table > Add

The following table describes the items in the previous figure.

Item	Description
Telephony OUI	The unique OUI that identifies the device manufacturer or vendor. The OUI is specified in three octet values (each octet is represented as two hexadecimal digits) separated by colons.
Description	Identifies the manufacturer or vendor associated with the OUI.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.7.2.3 OUI Based Auto VoIP

Use the OUI Based Auto VoIP page to configure the Organizationally Unique Identifier (OUI) based Auto VoIP priority and to enable or disable the Auto VoIP mode on the interfaces.

To access this page, click **QoS > Auto VoIP > OUI Based Auto VoIP**.

Figure 4.393 QoS > Auto VoIP > OUI Based Auto VoIP

The following table describes the items in the previous figure.

Item	Description
Auto VoIP VLAN	The VLAN used to segregate VoIP traffic from other non-voice traffic. All VoIP traffic that matches a value in the known OUI list gets assigned to this VoIP VLAN.

Item	Description
Priority	The 802.1p priority used for traffic that matches a value in the known OUI list. If the Auto VoIP mode is enabled and the interface detects an OUI match, the device assigns the traffic in that session to the traffic class mapped to this priority value. Traffic classes with a higher value are generally used for time-sensitive traffic.
Interface	The interface associated with the rest of the data in the row. When editing Auto VoIP settings on one or more interfaces, this field identifies the interface(s) being configured.
Auto VoIP Mode	The administrative mode of OUI-based Auto VoIP on the interface.
Operational Status	The operational status of an interface. To be up, an interface must be administratively enabled and have a link.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.
Edit All	Click Edit All to apply the same settings to all interfaces.
Cancel	Click Cancel to restore default value.

4.7.2.4 Protocol Based Auto VoIP

Use the Protocol Based Auto VoIP page to configure the protocol-based Auto VoIP priority settings and to enable or disable the protocol-based Auto VoIP mode on the interfaces.

To access this page, click **QoS > Auto VoIP > Protocol Based Auto VoIP**.

The screenshot shows the configuration page for Protocol Based Auto VoIP. At the top, there are configuration fields: 'Auto VoIP VLAN' (set to 'Not Configured'), 'Prioritization Type' (radio buttons for 'Remark' and 'Traffic Class'), '802.1p Priority' (input field with '0 to 7' range), and 'Traffic Class' (input field with '0 to 7' range). Below this is a table with 10 rows, one for each interface from 0/1 to 0/10. Each row has three columns: 'Interface', 'Auto VoIP Mode', and 'Operational Status'. All 'Auto VoIP Mode' values are 'Disable' and all 'Operational Status' values are 'Down'. At the bottom of the page are buttons for 'Submit', 'Refresh', 'Edit', 'Edit All', and 'Cancel'.

Figure 4.394 QoS > Auto VoIP > Protocol Based Auto VoIP

The following table describes the items in the previous figure.

Item	Description
Auto VoIP VLAN	The VLAN used to segregate VoIP traffic from other non-voice traffic. All VoIP traffic in a session identified by the call-control protocol gets assigned to this VoIP VLAN.
Prioritization Type	The method used to prioritize VoIP traffic when a call-control protocol is detected, which is one of the following: <ul style="list-style-type: none"> ■ Remark: Remark the voice traffic with the specified 802.1p priority value at the ingress interface. ■ Traffic Class: Assign VoIP traffic to the specified traffic class when egressing the interface.

Item	Description
802.1p Priority	The 802.1p priority used for protocol-based VoIP traffic. This field can be configured if the Prioritization Type is 802.1p Priority. If the Auto VoIP mode is enabled and the interface detects a call-control protocol, the device marks traffic in that session with the specified 802.1p priority value to ensure voice traffic always gets the highest priority throughout the network path. Egress tagging must be administratively enabled on the appropriate uplink port to carry the remarked priority at the egress port.
Traffic Class	The traffic class used for protocol-based VoIP traffic. This field can be configured if the Prioritization Type is Traffic Class. If the Auto VoIP mode is enabled and the interface detects a call-control protocol, the device assigns the traffic in that session to the configured Class of Service (CoS) queue. Traffic classes with a higher value are generally used for time-sensitive traffic. The CoS queue associated with the specified traffic class should be configured with the appropriate bandwidth allocation to allow priority treatment for VoIP traffic.
Interface	The interface associated with the rest of the data in the row. When editing Auto VoIP settings on one or more interfaces, this field identifies the interface(s) being configured.
Auto VoIP Mode	<p>The administrative mode of the Auto VoIP feature on the interface:</p> <ul style="list-style-type: none"> ■ Enable: The interface scans incoming traffic for the following call-control protocols: <ul style="list-style-type: none"> – Session Initiation Protocol (SIP) – H.323 – Skinny Client Control Protocol (SCCP) ■ Disable: The interface does not use the Auto VoIP feature to scan for call-control protocols.
Operational Status	The operational status of an interface. To be up, an interface must be administratively enabled and have a link.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.
Edit All	Click Edit All to apply the same settings to all interfaces.
Cancel	Click Cancel to restore default value.

4.7.3 Class of Service

4.7.3.1 IP DSCP

Use the CoS IP DSCP Mapping Configuration page to configure the per-interface mapping between the IP DiffServ Code Point (DSCP) value and the traffic class. A DSCP value can be included in the Service Type field of an IP header. When traffic is queued for transmission on the interface, the DSCP value in the IP header is mapped to the traffic class specified on this page. A traffic class with a higher value has priority over a traffic class with a lower value.

To access this page, click **QoS > Class of Service > IP DSCP**.

The screenshot shows a configuration page for IP DSCP mapping. At the top, there is a dropdown menu for 'Interface' set to 'Global'. Below this is a table with two columns: 'IP DSCP' and 'Traffic Class'. The 'IP DSCP' column lists values from 0 to 22. The 'Traffic Class' column contains radio buttons for values 0 through 7. In the first row (IP DSCP 0), the radio button for Traffic Class 1 is selected. This pattern repeats for other rows, with different Traffic Class values being selected for different IP DSCP values.

Figure 4.395 QoS > Class of Service > IP DSCP

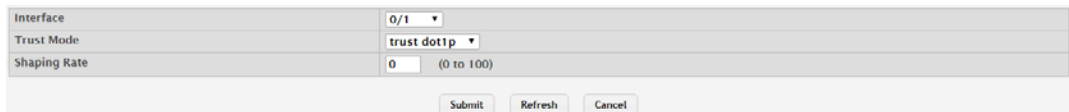
The following table describes the items in the previous figure.

Item	Description
Interface	The interface to configure. To configure the same IP DSCP-to-Traffic Class mappings on all interfaces, select the Global menu option.
IP DSCP	The list of possible IP DSCP values the IP header can include.
Traffic Class	The internal traffic class to which the corresponding IP DSCP priority value is mapped. The higher the traffic class value, the higher its priority is for sending traffic.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.

4.7.3.2 Interface

Use the CoS Interface Configuration page to configure the per-interface Class of Service (CoS) settings. The CoS feature allows preferential treatment for certain types of traffic over others. To set up this preferential treatment, you can configure the CoS interface settings and individual queues on the egress ports to provide customization that suits the network environment. The level of service is determined by the egress port queue to which the traffic is assigned. When traffic is queued for transmission, the rate at which it is serviced depends on how the queue is configured and possibly the amount of traffic present in other queues for that port.

To access this page, click **QoS > Class of Service > Interface**.



The screenshot shows a configuration form with three rows: 'Interface' with a dropdown menu showing '0/1', 'Trust Mode' with a dropdown menu showing 'trust dot1p', and 'Shaping Rate' with a text input field showing '0' and '(0 to 100)' next to it. Below the form are three buttons: 'Submit', 'Refresh', and 'Cancel'.

Figure 4.396 QoS > Class of Service > Interface

The following table describes the items in the previous figure.

Item	Description
Interface	The interface to configure. To configure the same settings on all interfaces, select the Global menu option.
Trust Mode	The trust mode for ingress traffic on the interface, which is one of the following: <ul style="list-style-type: none">■ untrusted: The interface ignores any priority designations encoded in incoming packets, and instead sends the packets to a traffic queue based on the ingress port's default priority.■ trust dot1p: The port accepts at face value the 802.1p priority designation encoded within packets arriving on the port.■ trust IP DSCP: The port accepts at face value the IP DSCP priority designation encoded within packets arriving on the port.
Shaping Rate	The upper limit on how much traffic can leave a port. The limit on maximum transmission bandwidth has the effect of smoothing temporary traffic bursts over time so that the transmitted traffic rate is bounded. The specified value represents a percentage of the maximum negotiated bandwidth.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.7.3.3 Queue

Use the CoS Interface Queue Configuration page to define the behavior of the egress CoS queues on each interface. User-configurable parameters control the amount of bandwidth used by the queue, the queue depth during times of congestion, and the scheduling of packet transmission from the set of all queues on an interface. Each interface has its own CoS queue-related configuration.

To access this page, click **QoS > Class of Service > Queue**.

Queue ID	Minimum Bandwidth (%)	Scheduler Type	Queue Management Type
0	0	Weighted	TailDrop
1	0	Weighted	TailDrop
2	0	Weighted	TailDrop
3	0	Weighted	TailDrop
4	0	Weighted	TailDrop
5	0	Weighted	TailDrop
6	0	Weighted	TailDrop
7	0	Weighted	TailDrop

Figure 4.397 QoS > Class of Service > Queue

The following table describes the items in the previous figure.

Item	Description
Interface	The interface to configure. To configure the same settings on all interfaces, select the Global menu option.
Total Minimum Bandwidth Allocation	Shows the total minimum bandwidth allocation to the selected interface for all the queues.
Queue ID	The CoS queue. The higher the queue value, the higher its priority is for sending traffic.
Minimum Bandwidth (%)	The minimum guaranteed bandwidth allocated to the selected queue on the interface. Setting this value higher than its corresponding Maximum Bandwidth automatically increases the maximum to the same value. A zero value (0) means no guaranteed minimum. The sum of individual Minimum Bandwidth values for all queues in the selected interface cannot exceed defined maximum 100.
Scheduler Type	The type of queue processing. Defining this value on a per-queue basis allows you to create the desired service characteristics for different types of traffic. The options are as follows: <ul style="list-style-type: none"> Weighted: Weighted round robin associates a weight to each queue. Strict: Strict priority services traffic with the highest priority on a queue first.
Queue Management Type	The type of queue depth management techniques used for all queues on this interface. The options are as follows: <ul style="list-style-type: none"> Taildrop: All packets on a queue are safe until congestion occurs. At this point, any additional packets queued are dropped.
Refresh	Click Refresh to update the screen.
Edit	Click Edit to edit the selected entries.
Restore Default	Click Restore Default to restore all CoS queue settings on the select interface to the default values.

4.7.4 Diffserv

4.7.4.1 Global

Use the Diffserv Global Configuration and Status page to configure the administrative mode of Differentiated Services (DiffServ) support on the device and to view the current and maximum number of entries in each of the main DiffServ private MIB tables. DiffServ allows traffic to be classified into streams and given certain QoS treatment in accordance with defined per-hop behaviors.

Packets are classified and processed based on defined criteria. The classification criteria is defined by a class. The processing is defined by a policy's attributes. Policy attributes may be defined on a per-class instance basis, and it is these attributes that are applied when a match occurs. A policy can contain multiples classes. When the policy is active, the actions taken depend on which class matches the packet.

To access this page, click **QoS > Diffserv > Global**.

MIB Table	Current Number / Maximum Number
Class Table	0 / 32
Class Rule Table	0 / 416
Policy Table	0 / 64
Policy Instance Table	0 / 1792
Policy Attribute Table	0 / 5376
Service Table	0 / 42

Figure 4.398 QoS > Diffserv > Global

The following table describes the items in the previous figure.

Item	Description
Diffserv Admin Mode	The administrative mode of DiffServ on the device. While disabled, the DiffServ configuration is retained and can be changed, but it is not active. While enabled, Differentiated Services are active.
MIB Table	
Class Table	The current and maximum number of classifier entries in the table. DiffServ classifiers differentiate among traffic types.
Class Rule Table	The current and maximum number of class rule entries in the table. Class rules specify the match criteria that belong to a class definition.
Policy Table	The current and maximum number of policy entries in the table. The policy determines the traffic conditioning or service provisioning actions applied to a traffic class.
Policy Instance Table	The current and maximum number of policy-class instance entries in the table. A policy-class instance is a policy that is associated with an existing DiffServ class.
Policy Attribute Table	The current and maximum number of policy attribute entries in the table. A policy attribute entry attaches various policy attributes to a policy-class instance.
Service Table	The current and maximum number of service entries in the table. A service entry associates a DiffServ policy with an interface and inbound or outbound direction.
Submit	Click Submit to save the values and update the screen.
Refresh	Click Refresh to update the screen.
Cancel	Click Cancel to restore default value.

4.7.4.2 Class Summary

Use the Diffserv Class Summary page to create or remove DiffServ classes and to view summary information about the classes that exist on the device. Creating a class is the first step in using DiffServ to provide Quality of Service. After a class is created, you can define the match criteria for the class.

To access this page, click **QoS > Diffserv > Class Summary**.




Figure 4.399 QoS > Diffserv > Class Summary

The following table describes the items in the previous figure.

Item	Description
Name	The name of the DiffServ class. When adding a new class or renaming an existing class, the name of the class is specified in the Class field of the dialog window.
Type	The class type, which is one of the following: <ul style="list-style-type: none"> ■ All: All the various match criteria defined for the class should be satisfied for a packet match. All signifies the logical AND of all the match criteria.
Protocol	The Layer 3 protocol to use for filtering class types, which is either IPv4 or IPv6.
Match Criteria	The criteria used to match packets.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new DiffServ class. See the following procedure.
Rename	Click Rename to rename the name of an existing class.
Remove	Click Remove to remove the selected entries.

To add a new DiffServ class:

Click **QoS > Diffserv > Class Summary > Add**.



Figure 4.400 QoS > Diffserv > Class Summary > Add

The following table describes the items in the previous figure.

Item	Description
Class	Enter the name of the DiffServ class.
Type	The class type, which is one of the following: <ul style="list-style-type: none"> ■ All: All the various match criteria defined for the class should be satisfied for a packet match. All signifies the logical AND of all the match criteria.
Protocol	The Layer 3 protocol to use for filtering class types, which is either IPv4 or IPv6.

Item	Description
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.7.4.3 Class Configuration

Use the Diffserv Class Configuration page to define the criteria to associate with a DiffServ class. As packets are received or transmitted, these DiffServ classes are used to classify and prioritize packets. Each class can contain multiple match criteria. To access this page, click **QoS > Diffserv > Class Configuration**.

Figure 4.401 QoS > Diffserv > Class Configuration

The following table describes the items in the previous figure.

Item	Description
Class	The name of the class. To configure match criteria for a class, select its name from the menu.
Type	The class type, which is one of the following: <ul style="list-style-type: none"> ■ All: All the various match criteria defined for the class should be satisfied for a packet match. All signifies the logical AND of all the match criteria.
L3 Protocol	The Layer 3 protocol to use for filtering class types, which is either IPv4 or IPv6.
Match Criteria	The type of match criteria defined for the selected class.
Value	The configured value of the match criteria that corresponds to the match type.
Refresh	Click Refresh to update the screen.
Add Match Criteria	Click Add Match Criteria to define criteria for matching packets within a class. See the following procedure.
Remove Reference Class	Click Remove Reference Class to remove the associated reference class from the selected class.

To define criteria for matching packets within a class:
 Click **QoS > Diffserv > Class Configuration > Add Match Criteria**.



Figure 4.402 QoS > Diffserv > Class Configuration > Add Match Criteria

The following table describes the items in the previous figure.

Item	Description
Any	Select this option to specify that all packets are considered to match the specified class. There is no need to configure additional match criteria if Any is selected because a match will occur on all packets.
Reference Class	Select this option to reference another class for criteria. The match criteria defined in the referenced class is as match criteria in addition to the match criteria you define for the selected class. After selecting this option, the classes that can be referenced are displayed. Select the class to reference. A class can reference at most one other class of the same type.
CoS	Select this option to require the Class of Service (CoS) value in an Ethernet frame header to match the specified CoS value.
Secondary CoS	Select this option to require the secondary CoS value in an Ethernet frame header to match the specified secondary CoS value.
Ethertype	Select this option to require the EtherType value in the Ethernet frame header to match the specified EtherType value. After you select this option, specify the EtherType value in one of the following two fields: <ul style="list-style-type: none"> ■ EtherType Keyword: The menu includes several common protocols that are mapped to their EtherType values. ■ EtherType Value: This field accepts custom EtherType values.
VLAN	Select this option to require a packet's VLAN ID to match a VLAN ID or a VLAN ID within a continuous range. If you configure a range, a match occurs if a packet's VLAN ID is the same as any VLAN ID within the range. After you select this option, use the following fields to configure the VLAN match criteria: <ul style="list-style-type: none"> ■ VLAN ID: The VLAN ID to match.

Item	Description
Secondary VLAN	<p>Select this option to require a packet's VLAN ID to match a secondary VLAN ID or a secondary VLAN ID within a continuous range. If you configure a range, a match occurs if a packet's secondary VLAN ID is the same as any secondary VLAN ID within the range. After you select this option, use the following fields to configure the secondary VLAN match criteria:</p> <ul style="list-style-type: none"> ■ Secondary VLAN ID: The secondary VLAN ID to match.
Source MAC Address	<p>Select this option to require a packet's source MAC address to match the specified MAC address. After you select this option, use the following fields to configure the source MAC address match criteria:</p> <ul style="list-style-type: none"> ■ MAC Address: The source MAC address to match. ■ MAC Mask: The MAC mask, which specifies the bits in the source MAC address to compare against an Ethernet frame. Use F's and zeros to configure the MAC mask. An F means that the bit is checked, and a zero in a bit position means that the data is not significant. For example, if the MAC address is aa:bb:cc:dd:ee:ff, and the mask is ff:ff:00:00:00:00, all MAC addresses with aa:bb:xx:xx:xx:xx result in a match (where x is any hexadecimal number). Note that this is not a wildcard mask, which ACLs use.
Destination MAC Address	<p>Select this option to require a packet's destination MAC address to match the specified MAC address. After you select this option, use the following fields to configure the destination MAC address match criteria:</p> <ul style="list-style-type: none"> ■ MAC Address: The destination MAC address to match. ■ MAC Mask: The MAC mask, which specifies the bits in the destination MAC address to compare against an Ethernet frame. Use F's and zeros to configure the MAC mask. An F means that the bit is checked, and a zero in a bit position means that the data is not significant. For example, if the MAC address is aa:bb:cc:dd:ee:ff, and the mask is ff:ff:00:00:00:00, all MAC addresses with aa:bb:xx:xx:xx:xx result in a match (where x is any hexadecimal number). Note that this is not a wildcard mask, which ACLs use.
Source IP Address	<p>Select this option to require the source IP address in a packet header to match the specified values. After you select this option, use the following fields to configure the source IP address match criteria:</p> <ul style="list-style-type: none"> ■ IP Address: The source IP address to match. ■ IP Mask: A valid subnet mask, which determines the bits in the IP address that are significant. Note that this is not a wildcard mask.
Destination IP Address	<p>Select this option to require the destination IP address in a packet header to match the specified values. After you select this option, use the following fields to configure the destination IP address match criteria:</p> <ul style="list-style-type: none"> ■ IP Address: The destination IP address to match. ■ IP Mask: A valid subnet mask, which determines the bits in the IP address that are significant. Note that this is not a wildcard mask.
Source IPv6 Address	<p>Select this option to require the source IPv6 address in a packet header to match the specified values. After you select this option, use the following fields to configure the source IPv6 address match criteria:</p> <ul style="list-style-type: none"> ■ Source Prefix: The source IPv6 prefix to match. ■ Source Prefix Length: The IPv6 prefix length.

Item	Description
Destination IPv6 Address	<p>Select this option to require the destination IPv6 address in a packet header to match the specified values. After you select this option, use the following fields to configure the destination IPv6 address match criteria:</p> <ul style="list-style-type: none"> ■ Destination Prefix: The destination IPv6 prefix to match. ■ Destination Prefix Length: The IPv6 prefix length.
Source L4 Port	<p>Select this option to require a packet's TCP/UDP source port to match the specified port or the port number within a range of port numbers. If you configure a range, a match occurs if a packet's source port number is the same as any source port number within the range. After you select this option, use the following fields to configure a source port keyword, source port number, or source port range for the match criteria:</p> <ul style="list-style-type: none"> ■ Protocol: Select the desired L4 keyword from the list on which the match is based. If you select a keyword, the other source port configuration fields are not configurable. ■ Port: The source port number to match.
Destination L4 Port	<p>Select this option to require a packet's TCP/UDP destination port to match the specified port or the port number within a range of port numbers. If you configure a range, a match occurs if a packet's destination port number is the same as any destination port number within the range. After you select this option, use the following fields to configure a destination port keyword, destination port number, or destination port range for the match criteria:</p> <ul style="list-style-type: none"> ■ Protocol: Select the desired L4 keyword from the list on which the match is based. If you select a keyword, the other destination port configuration fields are not configurable. ■ Port: The destination port number to match.
IP DSCP	<p>Select this option to require the packet's IP DiffServ Code Point (DSCP) value to match the specified value. The DSCP value is defined as the high-order six bits of the Service Type octet in the IP header. After you select this option, use one of the following fields to configure the IP DSCP match criteria:</p> <ul style="list-style-type: none"> ■ IP DSCP Keyword: The IP DSCP keyword code that corresponds to the IP DSCP value to match. If you select a keyword, you cannot configure an IP DSCP Value. ■ IP DSCP Value: The IP DSCP value to match.
IP Precedence	<p>Select this option to require the packet's IP Precedence value to match the number configured in the IP Precedence Value field. The IP Precedence field in a packet is defined as the high-order three bits of the Service Type octet in the IP header.</p>
IP TOS	<p>Select this option to require the packet's Type of Service (ToS) bits in the IP header to match the specified value. The IP ToS field in a packet is defined as all eight bits of the Service Type octet in the IP header. After you select this option, use the following fields to configure the ToS match criteria:</p> <ul style="list-style-type: none"> ■ IP TOS Bits: Enter a two-digit hexadecimal number to match the bits in a packet's ToS field. ■ IP TOS Mask: Specify the bit positions that are used for comparison against the IP ToS field in a packet.

Item	Description
Protocol	Select this option to require a packet header's Layer 4 protocol to match the specified value. After you select this option, use one of the following fields to configure the protocol match criteria: <ul style="list-style-type: none"> ■ Protocol: The L4 keyword that corresponds to value of the IANA protocol number to match. If you select a keyword, you cannot configure a Protocol Value. ■ Protocol Value: The IANA L4 protocol number value to match.
Flow Label	Select this option to require an IPv6 packet's flow label to match the configured value. The flow label is a 20-bit number that is unique to an IPv6 packet, used by end stations to signify quality-of-service handling in routers.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.7.4.4 Policy Summary

Use the Diffserv Policy Summary page to create or remove DiffServ policies and to view summary information about the policies that exist on the device. A policy defines the QoS attributes for one or more traffic classes. A policy attribute identifies the action taken when a packet matches a class rule. A policy is applied to a packet when a class match within that policy is found.

To access this page, click **QoS > Diffserv > Policy Summary**.




Figure 4.403 QoS > Diffserv > Policy Summary

The following table describes the items in the previous figure.

Item	Description
Name	The name of the DiffServ policy. When adding a new policy or renaming an existing policy, the name of the policy is specified in the Policy field of the dialog window.
Type	The traffic flow direction to which the policy is applied: <ul style="list-style-type: none"> ■ In: The policy is specific to inbound traffic. ■ Out: The policy is specific to outbound traffic.
Member Classes	The DiffServ class or classes that have been added to the policy.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a new DiffServ policy. See the following procedure.
Rename	Click Rename to rename the name of an existing policy.
Remove	Click Remove to remove the selected entries.

To add a new DiffServ policy:

Click **QoS > Diffserv > Policy Summary > Add**.




Figure 4.404 QoS > Diffserv > Policy Summary > Add

The following table describes the items in the previous figure.

Item	Description
Policy	Enter the name of the policy.
Type	The traffic flow direction to which the policy is applied: <input type="checkbox"/> In: The policy is specific to inbound traffic. <input type="checkbox"/> Out: The policy is specific to outbound traffic.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.7.4.5 Policy Configuration

Use the DiffServ Policy Configuration page to add or remove a DiffServ policy-class association and to configure the policy attributes. The policy attributes identify the action or actions taken when a packet matches a class rule.

To access this page, click **QoS > Diffserv > Policy Configuration**.




Figure 4.405 QoS > Diffserv > Policy Configuration

The following table describes the items in the previous figure.

Item	Description
Policy	The name of the policy. To add a class to the policy, remove a class from the policy, or configure the policy attributes, you must first select its name from the menu.
Type	The traffic flow direction to which the policy is applied.
Class	The DiffServ class or classes associated with the policy. The policy is applied to a packet when a class match within that policy-class is found.
Policy Attribute Details	The policy attribute types and their associated values that are configured for the policy.
Refresh	Click Refresh to update the screen.
Add Class	Click Add Class to add a class to the policy. See the following procedure.
Add Attribute	Click Add Attribute to add attributes to a policy or to change the policy attributes. See the following procedure.
Remove Last Class	Click Remove Last Class to remove the most recently associated class from the selected policy.

To add a class to the policy:

Click **QoS > Diffserv > Policy Configuration > Add Class**.




Figure 4.406 QoS > Diffserv > Policy Configuration > Add Class

The following table describes the items in the previous figure.

Item	Description
Policy	The name of the policy. To add a class to the policy, remove a class from the policy, or configure the policy attributes, you must first select its name from the menu.
Type	The traffic flow direction to which the policy is applied.
Class	The DiffServ class or classes associated with the policy. The policy is applied to a packet when a class match within that policy-class is found.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

To add attributes to a policy or to change the policy attributes:

Click **QoS > Diffserv > Policy Configuration > Add Attribute**.



Figure 4.407 QoS > Diffserv > Policy Configuration > Add Attribute

The following table describes the items in the previous figure.

Item	Description
Assign Queue	Select this option to assign matching packets to a traffic queue. Use the Queue ID Value field to select the queue to which the packets of this policy-class are assigned.
Drop	Select this option to drop packets that match the policy-class.
Mark CoS	Select this option to mark all packets in a traffic stream with the specified Class of Service (CoS) queue value. Use the Class of Service field to select the CoS value to mark in the priority field of the 802.1p header (the only tag in a single tagged packet or the first or outer 802.1Q tag of a double VLAN tagged packet). If the packet does not already contain this header, one is inserted.
Mark CoS as Secondary CoS	Select this option to mark the priority field of the 802.1p header in the outer tag of a double-VLAN tagged packet with the same CoS value that is included in the inner tag.

Item	Description
Mark IP DSCP	<p>Select this option to mark all packets in the associated traffic stream with the specified IP DSCP value. After you select this option, use one of the following fields to configure the IP DSCP value to mark in packets that match the policy-class:</p> <ul style="list-style-type: none"> ■ IP DSCP Keyword: The IP DSCP keyword code that corresponds to the IP DSCP value. If you select a keyword, you cannot configure an IP DSCP Value. ■ IP DSCP Value: The IP DSCP value.
Mark IP Precedence	<p>Select this option to mark all packets in the associated traffic stream with the specified IP Precedence value. After you select this option, use the IP Precedence Value field to select the IP Precedence value to mark in packets that match the policy-class.</p>
Mirror Interface	<p>Select this option to copy the traffic stream to a specified egress port (physical or LAG) without bypassing normal packet forwarding. This action can occur in addition to any marking or policing action. It may also be specified along with a QoS queue assignment. Use the Interface menu to select the interface to which traffic is mirrored.</p>
Police Simple	<p>Select this option to enable the simple traffic policing style for the policy-class. The simple form of the police attribute uses a single data rate and burst size, resulting in two outcomes (conform and violate). After you select this option, configure the following policing criteria:</p> <ul style="list-style-type: none"> ■ Color Mode: The type of color policing used in DiffServ traffic conditioning. ■ Color Conform Class: For color-aware policing, packets in this class are metered against both the committed information rate (CIR) and the peak information rate (PIR). The class definition used for policing color awareness is only allowed to contain a single, non-excluded class match condition identifying one of the supported comparison fields: CoS, IP DSCP, IP Precedence, or Secondary COS. ■ Committed Rate (Kbps): The maximum allowed arrival rate of incoming packets for this class. ■ Committed Burst Size (Kbytes): The amount of conforming traffic allowed in a burst. ■ Conform Action: The action taken on packets that are considered conforming (below the police rate). ■ Violate Action: The action taken on packets that are considered non-conforming (above the police rate).

Item	Description
Police Single Rate	<p>Select this option to enable the single-rate traffic policing style for the policy-class. The single-rate form of the police attribute uses a single data rate and two burst sizes, resulting in three outcomes (conform, exceed, and violate). After you select this option, configure the following policing criteria:</p> <ul style="list-style-type: none"> ■ Color Mode: The type of color policing used in DiffServ traffic conditioning. ■ Color Conform Class: For color-aware policing, packets are metered against the committed information rate (CIR) and the peak information rate (PIR). The class definition used for policing color awareness is only allowed to contain a single, non-excluded class match condition identifying one of the supported comparison fields: CoS, IP DSCP, IP Precedence, or Secondary COS. This field is available only if one or more classes that meets the color-awareness criteria exist. ■ Color Exceed Class: For color-aware policing, packets are metered against the PIR only. ■ Committed Rate (Kbps): The maximum allowed arrival rate of incoming packets for this class. ■ Committed Burst Size (Kbytes): The amount of conforming traffic allowed in a burst. ■ Excess Burst Size (Kbytes): The amount of conforming traffic allowed to accumulate beyond the Committed Burst Size (Kbytes) value during longer-than-normal idle times. This value allows for occasional bursting. ■ Conform Action: The action taken on packets that are considered conforming (below the police rate). ■ Exceed Action: The action taken on packets that are considered to exceed the committed burst size but are within the excessive burst size. ■ Violate Action: The action taken on packets that are considered non-conforming (above the police rate).

Item	Description
Police Two Rate	<p>Select this option to enable the two-rate traffic policing style for the policy-class. The two-rate form of the police attribute uses two data rates and two burst sizes. Only the smaller of the two data rates is intended to be guaranteed. After you select this option, configure the following policing criteria:</p> <ul style="list-style-type: none"> ■ Color Mode: The type of color policing used in DiffServ traffic conditioning. ■ Color Conform Class: For color-aware policing, packets are metered against the committed information rate (CIR) and the peak information rate (PIR). The class definition used for policing color awareness is only allowed to contain a single, non-excluded class match condition identifying one of the supported comparison fields: CoS, IP DSCP, IP Precedence, or Secondary COS. This field is available only if one or more classes that meets the color-awareness criteria exist. ■ Color Exceed Class: For color-aware policing, packets are metered against the PIR. ■ Committed Rate (Kbps): The maximum allowed arrival rate of incoming packets for this class. ■ Committed Burst Size (Kbytes): The amount of conforming traffic allowed in a burst. ■ Peak Rate (Kbps): The maximum peak information rate for the arrival of incoming packets for this class. ■ Excess Burst Size (Kbytes): The maximum size of the packet burst that can be accepted to maintain the Peak Rate (Kbps). ■ Conform Action: The action taken on packets that are considered conforming (below the police rate). ■ Exceed Action: The action taken on packets that are considered to exceed the committed burst size but are within the excessive burst size. ■ Violate Action: The action taken on packets that are considered non-conforming (above the police rate).
Redirect Interface	Select this option to force a classified traffic stream to the specified egress port (physical port or LAG). Use the Interface field to select the interface to which traffic is redirected.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.7.4.6 Service Summary

Use the DiffServ Service Summary page to add DiffServ policies to interfaces, remove policies from interfaces, and edit policy-interface mappings.

To access this page, click **QoS > Diffserv > Service Summary**.




Figure 4.408 QoS > Diffserv > Service Summary

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. Only interfaces that have an associated policy are listed in the table.

Item	Description
Direction	The traffic flow direction to which the policy is applied: <ul style="list-style-type: none"> ■ Inbound: The policy is applied to traffic as it enters the interface. ■ Outbound: The policy is applied to traffic as it exits the interface.
Status	The status of the policy on the interface. A policy is Up if DiffServ is globally enabled, and if the interface is administratively enabled and has a link. Otherwise, the status is Down.
Policy	The DiffServ policy associated with the interface.
Refresh	Click Refresh to update the screen.
Add	Click Add to add a policy to an interface. See the following procedure.
Edit	Click Edit to edit the selected entries.
Remove	Click Remove to remove the selected entries.

To add a policy to an interface:

Click **QoS > Diffserv > Service Summary > Add**.

Figure 4.409 QoS > Diffserv > Service Summary > Add

The following table describes the items in the previous figure.

Item	Description
Interface	Select an interface to associate with a policy.
Policy In	The menu lists all policies configured with a type of In. Select the policy to apply to traffic as it enters the interface.
Policy Out	The menu lists all policies configured with a type of Out. Select the policy to apply to traffic as it exits the interface.
Submit	Click Submit to save the values.
Cancel	Click Cancel to close the window.

4.7.4.7 Service Statistics

The Diffserv Service Performance Statistics page displays service-level statistical information for all interfaces in the system to which a DiffServ policy has been attached.

To access this page, click **QoS > Diffserv > Service Statistics**.

Figure 4.410 QoS > Diffserv > Service Statistics

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. The table displays all interfaces that have a DiffServ policy currently attached in a traffic flow direction.
Direction	The traffic flow direction to which the policy is applied: <ul style="list-style-type: none"> ■ In: The policy is applied to traffic as it enters the interface. ■ Out: The policy is applied to traffic as it exits the interface.
Status	The operational status of this service interface, either Up or Down.
Refresh	Click Refresh to update the screen.

4.7.4.8 Policy Statistics

The Diffserv Policy Performance Statistics page displays class-oriented statistical information for the policy, which is specified by the interface and direction.

To access this page, click **QoS > Diffserv > Policy Statistics**.




Figure 4.411 QoS > Diffserv > Policy Statistics

The following table describes the items in the previous figure.

Item	Description
Interface	The interface associated with the rest of the data in the row. The table displays all interfaces that have a DiffServ policy currently attached in a traffic flow direction.
Direction	The traffic flow direction to which the policy is applied: <ul style="list-style-type: none"> ■ In: The policy is applied to traffic as it enters the interface. ■ Out: The policy is applied to traffic as it exits the interface.
Policy	The name of the policy currently attached to the interface.
Status	The operational status of the policy currently attached to the interface.
Class	The DiffServ class currently defined for the attached policy.
Packets Offered	The total number of packets offered to all class instances in this service policy before their defined DiffServ treatment is applied. This is the overall count per-interface, per-direction.
Packets Discarded	The total number of packets discarded for all class instances in this service policy for any reason due to DiffServ treatment. This is the overall count per-interface, per-direction.
Refresh	Click Refresh to update the screen.

Chapter 5

Command Line
Interface

5.1 Virtual Router Redundancy Protocol Commands

This section describes the commands you use to view and configure Virtual Router Redundancy Protocol (VRRP) and to view VRRP status information. VRRP helps provide failover and load balancing when you configure two devices as a VRRP pair.

5.1.1 Virtual Router Redundancy Protocol Commands

5.1.1.1 ip vrrp (Global Config)

Use this command in Global Config mode to enable the administrative mode of VRRP on the router.

Use the no command in Global Config mode to disable the default administrative mode of VRRP on the router.

```
ip vrrp
no ip vrrp
```

Default

None.

Command Mode

Global Config

5.1.1.2 ip vrrp (Interface Config)

Use this command in Interface Config mode to create a virtual router associated with the interface or range of interfaces. The parameter *vrid* is the virtual router ID, which has an integer value range from 1 to 255.

Use the no command in Interface Config mode to delete the virtual router associated with the interface. The virtual Router ID, *vrid*, is an integer value that ranges from 1 to 255.

```
ip vrrp vrid
no ip vrrp vrid
```

Command Mode

Interface Config

5.1.1.3 ip vrrp mode

This command enables the virtual router configured on the specified interface. Enabling the status field starts a virtual router. The parameter *vrid* is the virtual router ID which has an integer value ranging from 1 to 255.

Use the no command to disable the virtual router configured on the specified interface. Disabling the status field stops a virtual router.

```
ip vrrp vrid mode
no ip vrrp vrid mode
```

Command Mode

Interface Config

5.1.1.4 ip vrrp ip

This command sets the virtual router IP address value for an interface or range of interfaces. The value for *ipaddr* is the IP address which is to be configured on that interface for VRRP. The parameter *vrid* is the virtual router ID which has an integer value range from 1 to 255. You can use the optional [*secondary*] parameter to designate the IP address as a secondary IP address.

Use the no command in Interface Config mode to delete a secondary IP address value from the interface. To delete the primary IP address, you must delete the virtual router on the interface.

```
ip vrrp vrid ip ipaddr [secondary]
no ip vrrp vrid ipaddr secondary
```

Default

None.

Command Mode

Interface Config

5.1.1.5 ip vrrp accept-mode

Use this command to allow the VRRP Master to accept ping packets sent to one of the virtual router's IP addresses.

Use the no command to prevent the VRRP Master from accepting ping packets sent to one of the virtual router's IP addresses.

Note! *VRRP accept-mode allows only ICMP Echo Request packets. No other type of packet is allowed to be delivered to a VRRP address.*



```
ip vrrp vrid accept-mode
no ip vrrp vrid accept-mode
```

Default

Disabled.

Command Mode

Interface Config

5.1.1.6 ip vrrp authentication

This command sets the authorization details value for the virtual router configured on a specified interface or range of interfaces. The parameter {*none* | *simple*} specifies the authorization type for virtual router configured on the specified interface. The parameter [*key*] is optional, it is only required when authorization type is simple text password. The parameter *vrid* is the virtual router ID which has an integer value ranges from 1 to 255.

Use the no command to sets the default authorization details value for the virtual router configured on a specified interface or range of interfaces.

```
ip vrrp vrid authentication {none | simple key}
no ip vrrp vrid authentication
```

Default

No authorization.

Command Mode

Interface Config

5.1.1.7 **ip vrrp preempt**

This command sets the preemption mode value for the virtual router configured on a specified interface or range of interfaces. The parameter *vrid* is the virtual router ID, which is an integer from 1 to 255.

Use the no command to set the default preemption mode value for the virtual router configured on a specified interface or range of interfaces.

```
ip vrrp vrid preempt
no ip vrrp vrid preempt
```

Default

Enabled.

Command Mode

Interface Config

5.1.1.8 **ip vrrp priority**

This command sets the priority of a router within a VRRP group. It can be used to configure an interface or a range of interfaces. Higher values equal higher priority. The range is from 1 to 254. The parameter *vrid* is the virtual router ID, whose range is from 1 to 255.

The router with the highest priority is elected master. If a router is configured with the address used as the address of the virtual router, the router is called the “address owner.” The priority of the address owner is always 255 so that the address owner is always master. If the master has a priority less than 255 (it is not the address owner) and you configure the priority of another router in the group higher than the master’s priority, the router will take over as master only if preempt mode is enabled.

Use the no command to set the default priority value for the virtual router configured on a specified interface or range of interfaces.

```
ip vrrp vrid priority 1-254
no ip vrrp vrid priority
```

Default

100 unless the router is the address owner, in which case its priority is automatically set to 255.

Command Mode

Interface Config

5.1.1.9 **ip vrrp timers advertise**

This command sets the frequency, in seconds, that an interface or range of interfaces on the specified virtual router sends a virtual router advertisement.

Use the no command to set the default virtual router advertisement value for an interface or range of interfaces.

```
ip vrrp vrid timers advertise 1-255
no ip vrrp vrid timers advertise
```

Default

1

Command Mode

Interface Config

5.1.1.10 ip vrrp track interface

Use this command to alter the priority of the VRRP router based on the availability of its interfaces. This command is useful for tracking interfaces that are not configured for VRRP. Only IP interfaces are tracked. A tracked interface is up if the IP on that interface is up. Otherwise, the tracked interface is down. You can use this command to configure a single interface or range of interfaces. The argument `unit/slot/port` corresponds to a physical routing interface or VLAN routing interface. The keyword `vlan` is used to specify the VLAN ID of the routing VLAN directly instead of in a `unit/slot/port` format.

When the tracked interface is down or the interface has been removed from the router, the priority of the VRRP router will be decremented by the value specified in the `priority` argument. When the interface is up for IP protocol, the priority will be incremented by the `priority` value.

A VRRP configured interface can track more than one interface. When a tracked interface goes down, then the priority of the router will be decreased by 10 (the default priority decrement) for each downed interface. The default priority decrement is changed using the `priority` argument. The default priority of the virtual router is 100, and the default decrement priority is 10. By default, no interfaces are tracked. If you specify just the interface to be tracked, without giving the optional priority, then the default priority will be set. The default priority decrement is 10.

Use the `no` command to remove the interface or range of interfaces from the tracked list or to restore the priority decrement to its default.

```
ip vrrp vrid track interface {unit/slot/port | vlan 1-4093}
[decrement priority]
no ip vrrp vrid track interface {unit/slot/port | vlan 1-
4093} [decrement]
```

Default

Priority: 10.

Command Mode

Interface Config

5.1.1.11 show ip vrrp interface stats

This command displays the statistical information about each virtual router configured on the switch. The argument `unit/slot/port` corresponds to a physical routing interface or VLAN routing interface. The keyword `vlan` is used to specify the VLAN ID of the routing VLAN directly instead of a `unit/slot/port` format.

```
show ip vrrp interface stats {unit/slot/port | vlan 1-4093}
vrid
```

Command Mode

- Privileged EXEC
- User EXEC

Display Parameters

Term	Definition
Uptime	The time that the virtual router has been up, in days, hours, minutes and seconds.
Protocol	The protocol configured on the interface.
State Transitioned to Master	The total number of times virtual router state has changed to MASTER.
Advertisement Received	The total number of VRRP advertisements received by this virtual router.

Term	Definition
Advertisement Interval Errors	The total number of VRRP advertisements received for which advertisement interval is different than the configured value for this virtual router.
Authentication Failure	The total number of VRRP packets received that don't pass the authentication check.
IP TTL errors	The total number of VRRP packets received by the virtual router with IP TTL (time to live) not equal to 255.
Zero Priority Packets Received	The total number of VRRP packets received by virtual router with a priority of '0'.
Zero Priority Packets Sent	The total number of VRRP packets sent by the virtual router with a priority of '0'.
Invalid Type Packets Received	The total number of VRRP packets received by the virtual router with invalid 'type' field.
Address List Errors	The total number of VRRP packets received for which address list does not match the locally configured list for the virtual router.
Invalid Authentication Type	The total number of VRRP packets received with unknown authentication type.
Authentication Type Mismatch	The total number of VRRP advertisements received for which 'auth type' not equal to locally configured one for this virtual router.
Packet Length Errors	The total number of VRRP packets received with packet length less than length of VRRP header.

5.1.1.12 show ip vrrp

This command displays whether VRRP functionality is enabled or disabled on the switch. It also displays some global parameters which are required for monitoring. This command takes no options.

```
show ip vrrp
```

Command Mode

- Privileged EXEC
- User EXEC

Display Parameters

Term	Definition
VRRP Admin Mode	The administrative mode for VRRP functionality on the switch.
Router Checksum Errors	The total number of VRRP packets received with an invalid VRRP checksum value.
Router Version Errors	The total number of VRRP packets received with Unknown or unsupported version number.
Router VRID Errors	The total number of VRRP packets received with invalid VRID for this virtual router.

5.1.1.13 show ip vrrp interface

This command displays all configuration information and VRRP router statistics of a virtual router configured on a specific interface. The argument unit/slot/port corresponds to a physical routing interface or VLAN routing interface. The keyword vlan is the VLAN ID of the routing VLAN instead of in a unit/slot/port format. Use the output of the command to verify the track interface and track IP route configurations.

```
show ip vrrp interface {unit/slot/port | vlan 1-4093} vrid
```

Command Mode

- Privileged EXEC
- User EXEC

Example

The following shows example CLI display output for the command.

```
(Route)#show ip vrrp interface <u/s/p> vrid
```

```
Primary IP Address.....1.1.1.5
VMAC Address.....00:00:5e:00:01:01
Authentication Type.....None
Priority.....80
Configured priority.....100
Advertisement Interval (secs).....1
Pre-empt Mode.....Enable
Administrative Mode.....Enable
Accept Mode.....Enable
State.....Initialized
```

```
Track Interface   State   DecrementPriority
-----
<1/0/1>           down    10
```

```
TrackRoute (pfx/len)   State   DecrementPriority
-----
10.10.10.1/255.255.255.0  down    10
```

Display Parameters

Term	Definition
IP Address	The configured IP address for the Virtual router.
VMAC address	The VMAC address of the specified router.
Authentication type	The authentication type for the specific virtual router.
Priority	The priority value for the specific virtual router, taking into account any priority decrements for tracked interfaces or routes.
Configured Priority	The priority configured through the <code>ip vrrp vrid priority 1-254</code> command.
Advertisement interval	The advertisement interval in seconds for the specific virtual router.
Pre-Empt Mode	The preemption mode configured on the specified virtual router.
Administrative Mode	The status (Enable or Disable) of the specific router.
Accept Mode	When enabled, the VRRP Master can accept ping packets sent to one of the virtual router's IP addresses.
State	The state (Master/backup) of the virtual router.

5.1.1.14 **show ip vrrp interface brief**

This command displays information about each virtual router configured on the switch. This command takes no options. It displays information about each virtual router.

```
show ip vrrp interface brief
```

Command Mode

- Privileged EXEC
- User EXEC

Display Parameters

Term	Definition
Interface	<i>unit/slot/port</i>
VRID	The router ID of the virtual router.
IP Address	The virtual router IP address.
Mode	Indicates whether the virtual router is enabled or disabled.
State	The state (Master/backup) of the virtual router.

5.2 **Open Shortest Path First Commands**

This section describes the commands you use to view and configure Open Shortest Path First (OSPF), which is a link-state routing protocol that you use to route traffic within a network. This section contains the following subsections:

- “General OSPF Commands” on page 916
- “OSPF Interface Commands” on page 935
- “IP Event Dampening Commands” on page 941
- “OSPFv2 Stub Router Commands” on page 946
- “OSPF Show Commands” on page 947

5.2.1 **General OSPF Commands**

5.2.1.1 **router ospf**

Use this command to enable OSPF routing in a specified virtual router and to enter Router OSPF mode. If no virtual router is specified, OSPF routing is enabled in the default router.

```
router ospf [vrf vrf-name]
```

Parameters

Parameter	Description
<i>vrf vrf-name</i>	The virtual router on which to enable OSPF routing.

Command Mode

Global Config

5.2.1.2 **enable (OSPF)**

This command resets the default administrative mode of OSPF in the router (active). Use the no command to set the administrative mode of OSPF in the router to inactive.

```
enable
no enable
```

Default

Enabled.

Command Mode

Router OSPF Config

5.2.1.3 **network area (OSPF)**

Use this command to enable OSPFv2 on an interface and set its area ID if the IP address of an interface is covered by this network command.

Use the no command to disable the OSPFv2 on a interface if the IP address of an interface was earlier covered.

```
network ip-address wildcard-mask area area-id
no network ip-address wildcard-mask area area-id
```

Default

Disabled.

Command Mode

Router OSPF Config

5.2.1.4 **1583compatibility**

This command enables OSPF 1583 compatibility.

Use the no command to disable OSPF 1583 compatibility..

Note! *1583 compatibility mode is enabled by default. If all OSPF routers in the routing domain are capable of operating according to RFC 2328, OSPF 1583 compatibility mode should be disabled.*



```
1583compatibility
no 1583compatibility
```

Default

Enabled.

Command Mode

Router OSPF Config

5.2.1.5 **area default-cost (OSPF)**

This command configures the default cost for the stub area. You must specify the area ID and an integer value between 1-16777215.

```
area areaid default-cost 1-16777215
```

Command Mode

Router OSPF Config

5.2.1.6 **area nssa (OSPF)**

This command configures the specified areaid to function as an NSSA.

Use the no command to disable nssa from the specified area id.

```
area areaid nssa
no area areaid nssa
```

Command Mode

Router OSPF Config

5.2.1.7 **area nssa default-info-originate (OSPF)**

This command configures the metric value and type for the default route advertised into the NSSA. The optional metric parameter specifies the metric of the default route and is to be in a range of 1-16777214. If no metric is specified, the default value is ****. The metric type can be comparable (nssa-external 1) or noncomparable (nssa-external 2).

Use the no command to disable the default route advertised into the NSSA.

```
area areaid nssa default-info-originate [metric] [{compara-
ble | non-comparable}]
no area areaid nssa default-info-originate [metric] [{com-
parable | non-comparable}]
```

Command Mode

Router OSPF Config

5.2.1.8 **area nssa no-redistribute (OSPF)**

This command configures the NSSA Area Border router (ABR) so that learned external routes will not be redistributed to the NSSA.

Use the no command to disables the NSSA ABR so that learned external routes are redistributed to the NSSA.

```
area areaid nssa no-redistribute
no area areaid nssa no-redistribute
```

Command Mode

Router OSPF Config

5.2.1.9 **area nssa no-summary (OSPF)**

This command configures the NSSA so that summary LSAs are not advertised into the NSSA.

Use the no command to disables nssa from the summary LSAs.

```
area areaid nssa no-summary
no area areaid nssa no-summary
```

Command Mode

Router OSPF Config

5.2.1.10 area nssa translator-role (OSPF)

This command configures the translator role of the NSSA. A value of `always` causes the router to assume the role of the translator the instant it becomes a border router and a value of `candidate` causes the router to participate in the translator election process when it attains border router status.

Use the `no` command to disable the nssa translator role from the specified area id.

```
area areaid nssa translator-role {always | candidate}
no area areaid nssa translator-role {always | candidate}
```

Command Mode

Router OSPF Config

5.2.1.11 area nssa translator-stab-intv (OSPF)

This command configures the translator `stabilityinterval` of the NSSA. The `stabilityinterval` is the period of time that an elected translator continues to perform its duties after it determines that its translator status has been deposed by another router.

Use the `no` command to disable the nssa translator's `stabilityinterval` from the specified area id.

```
area areaid nssa translator-stab-intv stabilityinterval
no area areaid nssa translator-stab-intv stabilityinterval
```

Command Mode

Router OSPF Config

5.2.1.12 area range (OSPF)

Use the `area range` command in Router Configuration mode to configure a summary prefix that an area border router advertises for a specific area.

Use the `no` command to delete a specified area range or reverts an option to its default.

```
area areaid range prefix netmask {summarylink | nssaexternallink} [advertise | not-advertise] [cost cost]
no area areaid range prefix netmask {summarylink | nssaexternallink} [advertise | not-advertise] [cost cost]
```

Parameters

Parameter	Description
<code>areaid</code>	The area identifier for the area whose networks are to be summarized.
<code>prefix netmask</code>	The summary prefix to be advertised when the ABR computes a route to one or more networks within this prefix in this area.
<code>summarylink</code>	When this keyword is given, the area range is used when summarizing prefixes advertised in type 3 summary LSAs.
<code>nssaexternallink</code>	When this keyword is given, the area range is used when translating type 7 LSAs to type 5 LSAs.
<code>advertise</code>	(Optional) When this keyword is given, the summary prefix is advertised when the area range is active. This is the default.
<code>not-advertise</code>	[Optional] When this keyword is given, neither the summary prefix nor the contained prefixes are advertised when the area range is active. When the <code>not-advertise</code> option is given, any static cost previously configured is removed from the system configuration.

Parameter	Description
<code>cost cost</code>	[Optional] If an optional cost is given, OSPF sets the metric field in the summary LSA to the configured value rather than setting the metric to the largest cost among the networks covered by the area range. A static cost may only be configured if the area range is configured to advertise the summary. The range is 0 to 16,777,215. If the cost is set to 16,777,215 for type 3 summarization, a type 3 summary LSA is not advertised, but contained networks are suppressed. This behavior is equivalent to specifying the not-advertise option. If the range is configured for type 7 to type 5 translation, a type 5 LSA is sent if the metric is set to 16,777,215; however, other routers will not compute a route from a type 5 LSA with this metric.

Default

No area ranges are configured by default. No cost is configured by default.

Command Mode

OSPFv2 Router Config

Example

The following shows an example of the command.

```
!! Create area range
(Router)(Config-router)#area 1 range 10.0.0.0 255.0.0.0 summarylink
!! Delete area range
(Router)(Config-router)#no area 1 range 10.0.0.0 255.0.0.0 summarylink
```

The no form may be used to revert the [advertise | not-advertise] option to its default without deleting the area range. Deleting and recreating the area range would cause OSPF to temporarily advertise the prefixes contained within the range. Note that using either the advertise or not-advertise keyword reverts the configuration to the default. For example:

```
!! Create area range. Suppress summary.
(Router)(Config-router)#area 1 range 10.0.0.0 255.0.0.0 summarylink not-advertise
!! Advertise summary.
(Router)(Config-router)#no area 1 range 10.0.0.0 255.0.0.0 summarylink not-advertise
```

The no form may be use to remove a static area range cost, so that OSPF sets the cost to the largest cost among the contained routes.

```
!! Create area range with static cost.
(Router)(Config-router)#area 1 range 10.0.0.0 255.0.0.0 summarylink cost 1000
!! Remove static cost.
(Router)(Config-router)#no area 1 range 10.0.0.0 255.0.0.0 summarylink cost
```

5.2.1.13 area stub (OSPF)

This command creates a stub area for the specified area ID. A stub area is characterized by the fact that AS External LSAs are not propagated into the area. Removing AS External LSAs and Summary LSAs can significantly reduce the link state database of routers within the stub area.

Use the no command to delete a stub area for the specified area ID.

```
area areaid stub
no area areaid stub
```

Command Mode

Router OSPF Config

5.2.1.14 **area stub no-summary (OSPF)**

This command configures the Summary LSA mode for the stub area identified by *areaid*. Use this command to prevent LSA Summaries from being sent.

Use the no command to configure the default Summary LSA mode for the stub area identified by *areaid*.

```
area areaid stub no-summary  
no area areaid stub no-summary
```

Default

Disabled.

Command Mode

Router OSPF Config

5.2.1.15 **area virtual-link (OSPF)**

This command creates the OSPF virtual interface for the specified *areaid* and *neighbor*. The *neighbor* parameter is the Router ID of the neighbor.

Use the no command to delete the OSPF virtual interface from the given interface, identified by *areaid* and *neighbor*. The *neighbor* parameter is the Router ID of the neighbor.

```
area areaid virtual-link neighbor  
no area areaid virtual-link neighbor
```

Command Mode

Router OSPF Config

5.2.1.16 **area virtual-link authentication**

This command configures the authentication type and key for the OSPF virtual interface identified by *areaid* and *neighbor*. The *neighbor* parameter is the Router ID of the neighbor. The value for type is either none, simple, or encrypt. The *key* is composed of standard displayable, noncontrol keystrokes from a Standard 101/102-key keyboard. The authentication key must be 8 bytes or less if the authentication type is simple. If the type is encrypt, the key may be up to 16 bytes. Unauthenticated interfaces do not need an authentication key. If the type is encrypt, a key id in the range of 0 and 255 must be specified. The default value for authentication type is none. Neither the default password key nor the default key id are configured.

Use the no command to configures the default authentication type for the OSPF virtual interface identified by *areaid* and *neighbor*. The *neighbor* parameter is the Router ID of the neighbor.

```
area areaid virtual-link neighbor authentication {none |  
{simple key} | {encrypt key keyid}}  
no area areaid virtual-link neighbor authentication
```

Default

None.

Command Mode

Router OSPF Config

5.2.1.17 **area virtual-link dead-interval (OSPF)**

This command configures the dead interval for the OSPF virtual interface on the virtual interface identified by *areaid* and *neighbor*. The *neighbor* parameter is the Router ID of the neighbor. The range for seconds is 1 to 65535.

Use the no command to configure the default dead interval for the OSPF virtual interface on the virtual interface identified by *areaid* and *neighbor*. The *neighbor* parameter is the Router ID of the neighbor.

```
area areaid virtual-link neighbor dead-interval seconds  
no area areaid virtual-link neighbor dead-interval
```

Default

40

Command Mode

Router OSPF Config

5.2.1.18 **area virtual-link hello-interval (OSPF)**

This command configures the hello interval for the OSPF virtual interface on the virtual interface identified by *areaid* and *neighbor*. The *neighbor* parameter is the Router ID of the neighbor. The range for seconds is 1 to 65535.

Use the no command to configures the default hello interval for the OSPF virtual interface on the virtual interface identified by *areaid* and *neighbor*. The *neighbor* parameter is the Router ID of the neighbor.

```
area areaid virtual-link neighbor hello-interval 1-65535  
no areaid virtual-link neighbor hello-interval
```

Default

10

Command Mode

Router OSPF Config

5.2.1.19 **area virtual-link retransmit-interval (OSPF)**

This command configures the retransmit interval for the OSPF virtual interface on the virtual interface identified by *areaid* and *neighbor*. The *neighbor* parameter is the Router ID of the neighbor. The range for seconds is 0 to 3600.

Use the no command to configure the default retransmit interval for the OSPF virtual interface on the virtual interface identified by *areaid* and *neighbor*. The *neighbor* parameter is the Router ID of the neighbor.

```
area areaid virtual-link neighbor retransmit-interval seconds  
no area areaid virtual-link neighbor retransmit-interval
```

Default

5

Command Mode

Router OSPF Config

5.2.1.20 area virtual-link transmit-delay (OSPF)

This command configures the transmit delay for the OSPF virtual interface on the virtual interface identified by *areaid* and *neighbor*. The *neighbor* parameter is the Router ID of the neighbor. The range for seconds is 0 to 3600 (1 hour).

Use the no command to reset the default transmit delay for the OSPF virtual interface to the default value.

```
area areaid virtual-link neighbor transmit-delay seconds  
no area areaid virtual-link neighbor transmit-delay
```

Default

1

Command Mode

Router OSPF Config

5.2.1.21 auto-cost (OSPF)

By default, OSPF computes the link cost of each interface from the interface bandwidth. Faster links have lower metrics, making them more attractive in route selection. The configuration parameters in the *auto-cost* *reference bandwidth* and *bandwidth* commands give you control over the default link cost. You can configure for OSPF an interface bandwidth that is independent of the actual link speed. A second configuration parameter allows you to control the ratio of interface bandwidth to link cost. The link cost is computed as the ratio of a reference bandwidth to the interface bandwidth ($\text{ref_bw} / \text{interface bandwidth}$), where interface bandwidth is defined by the *bandwidth* command. Because the default reference bandwidth is 100 Mbps, OSPF uses the same default link cost for all interfaces whose bandwidth is 100 Mbps or greater. Use the *auto-cost* command to change the reference bandwidth, specifying the reference bandwidth in megabits per second (Mbps). The reference bandwidth range is 1-4294967 Mbps.

Use the no command to set the reference bandwidth to the default value.

```
auto-cost reference-bandwidth 1-4294967  
no auto-cost reference-bandwidth
```

Default

100 Mbps.

Command Mode

Router OSPF Config

5.2.1.22 **capability opaque**

Use this command to enable Opaque Capability on the Router. The information contained in Opaque LSAs may be used directly by OSPF or indirectly by an application wishing to distribute information throughout the OSPF domain. FASTPATH supports the storing and flooding of Opaque LSAs of different scopes. The default value of `enabled` means that OSPF will forward opaque LSAs by default. If you want to upgrade from a previous release, where the default was disabled, opaque LSA forwarding will be enabled. If you want to disable opaque LSA forwarding, then you should enter the command `no capability opaque` in OSPF Router Configuration mode after the software upgrade.

Use the `no` command to disable opaque capability on the router.

```
capability opaque
no capability opaque
```

Default

Enabled.

Command Mode

Router Config

5.2.1.23 **clear ip ospf**

Use this command to disable and re-enable OSPF for the specified virtual router. If no virtual router is specified, the default router is disabled and re-enabled.

```
clear ip ospf [vrf vrf-name]
```

Command Mode

Privileged EXEC

5.2.1.24 **clear ip ospf configuration**

Use this command to reset the OSPF Configuration to factory defaults for the specified virtual router. If no virtual router is specified, the default router is cleared.

```
clear ip ospf configuration [vrf vrf-name]
```

Command Mode

Privileged EXEC

5.2.1.25 **clear ip ospf counters**

Use this command to reset global and interface statistics for the specified virtual router. If no virtual router is specified, the global and interface statistics are reset for the default router.

```
clear ip ospf counters
```

Command Mode

Privileged EXEC

5.2.1.26 **clear ip ospf neighbor**

Use this command to drop the adjacency with all OSPF neighbors for the specified virtual router. On each neighbor's interface, send a one-way hello. Adjacencies may then be re-established. If no router is specified, adjacency with all OSPF neighbors is dropped for the default router. To drop all adjacencies with a specific router ID, specify the neighbor's Router ID using the optional parameter [`neighbor-id`].

```
clear ip ospf neighbor [vrf vrf-name] [neighbor-id]
```

Command Mode

Privileged EXEC

5.2.1.27 clear ip ospf neighbor interface

To drop adjacency with all neighbors on a specific interface, use the optional parameter `[unit/slot/port]`. To drop adjacency with a specific router ID on a specific interface, use the optional parameter `[neighbor-id]`.

```
clear ip ospf neighbor interface [unit/slot/port] [neighbor-id]
```

Command Mode

Privileged EXEC

5.2.1.28 clear ip ospf redistribution

Use this command to flush all self-originated external LSAs for the specified virtual router. If no router is specified, the command is executed for the default router. Reapply the redistribution configuration and reoriginate prefixes as necessary.

```
clear ip ospf redistribution [vrf vrf-name]
```

Command Mode

Privileged EXEC

5.2.1.29 default-information originate (OSPF)

This command is used to control the advertisement of default routes.

Use the no command to control the advertisement of default routes.

```
default-information originate [always] [metric 0-16777214]
[metric-type {1 | 2}]
no default-information originate [metric] [metric-type]
```

Default

- metric - unspecified
- type - 2

Command Mode

Router OSPF Config

5.2.1.30 default-metric (OSPF)

This command is used to set a default for the metric of distributed routes.

Use the no command to set a default for the metric of distributed routes.

```
default-metric 1-16777214
no default-metric
```

Command Mode

Router OSPF Config

5.2.1.31 **distance ospf (OSPF)**

This command sets the route preference value of OSPF in the router. Lower route preference values are preferred when determining the best route. The type of OSPF route can be *intra*, *inter*, or *external*. All the external type routes are given the same preference value. The range of *preference* value is 1 to 255.

Use the *no* command to set the default route preference value of OSPF routes in the router. The type of OSPF can be *intra*, *inter*, or *external*. All the external type routes are given the same preference value.

```
distance ospf {intra-area 1-255 | inter-area 1-255 | external 1-255}
no distance ospf {intra-area | inter-area | external}
```

Default

110

Command Mode

Router OSPF Config

5.2.1.32 **distribute-list out (OSPF)**

Use this command to specify the access list to filter routes received from the source protocol.

Use the *no* command to specify the access list to filter routes received from the source protocol.

```
distribute-list 1-199 out {rip | bgp | static | connected}
no distribute-list 1-199 out {rip | bgp | static | connected}
```

Command Mode

Router OSPF Config

5.2.1.33 **exit-overflow-interval (OSPF)**

This command configures the exit overflow interval for OSPF. It describes the number of seconds after entering overflow state that a router will wait before attempting to leave the overflow state. This allows the router to again originate nondefault AS-external-LSAs. When set to 0, the router will not leave overflow state until restarted. The range for seconds is 0 to 2147483647 seconds.

Use the *no* command to configure the default exit overflow interval for OSPF.

```
exit-overflow-interval seconds
no exit-overflow-interval
```

Default

0

Command Mode

Router OSPF Config

5.2.1.34 external-lsdb-limit (OSPF)

This command configures the external LSDB limit for OSPF. If the value is -1, then there is no limit. When the number of nondefault AS-external-LSAs in a router's link-state database reaches the external LSDB limit, the router enters overflow state. The router never holds more than the external LSDB limit nondefault AS-external-LSAs in its database. The external LSDB limit MUST be set identically in all routers attached to the OSPF backbone and/or any regular OSPF area. The range for limit is -1 to 2147483647.

Use the no command to configure the default external LSDB limit for OSPF.

```
external-lsdb-limit limit
no external-lsdb-limit
```

Default

-1

Command Mode

Router OSPF Config

5.2.1.35 log-adjacency-changes

To enable logging of OSPFv2 neighbor state changes, use the log-adjacency-changes command in Router Configuration mode. State changes are logged with INFORMATIONAL severity.

Use the no command to disable state change logging.

```
log-adjacency-changes [detail]
no log-adjacency-changes [detail]
```

Parameters

Parameter	Description
<code>detail</code>	(Optional) When this keyword is specified, all adjacency state changes are logged. Otherwise, OSPF only logs transitions to FULL state and when a backwards transition occurs.

Default

Adjacency state changes are logged, but without the detail option.

Command Mode

OSPFv2 Router Config

5.2.1.36 prefix-suppression (Router OSPF Config)

This command suppresses the advertisement of all the IPv4 prefixes except for prefixes that are associated with secondary IPv4 addresses, loopbacks, and passive interfaces from the OSPFv2 router advertisements.

To suppress a loopback or passive interface, use the `ip ospf prefix-suppression` command in Interface Configuration mode. Prefixes associated with secondary IPv4 addresses can never be suppressed.

Use the no command to disable prefix-suppression. No prefixes are suppressed from getting advertised.

```
prefix-suppression
no prefix-suppression
```

Default

Prefix suppression is disabled.

Command Mode

Router OSPF Config

5.2.1.37 **prefix-suppression (Router OSPFv3 Config)**

This command suppresses the advertisement of all the IPv6 prefixes except for prefixes that are associated with secondary IPv6 addresses, loopbacks, and passive interfaces from the OSPFv3 router advertisements.

To suppress a loopback or passive interface, use the `ipv ospf prefix-suppression` command in Interface Configuration mode. Prefixes associated with secondary IPv6 addresses can never be suppressed.

Use the `no` command to remove.

```
prefix-suppression
no prefix-suppression
```

Default

Prefix suppression is disabled.

Command Mode

Router OSPFv3 Config

5.2.1.38 **router-id (OSPF)**

This command sets a 4-digit dotted-decimal number uniquely identifying the router ospf id. The `ipaddress` is a configured value.

```
router-id ipaddress
```

Command Mode

Router OSPF Config

5.2.1.39 **redistribute (OSPF)**

This command configures OSPF protocol to allow redistribution of routes from the specified source protocol/routers.

Use the `no` command to configure OSPF protocol to prohibit redistribution of routes from the specified source protocol/routers.

```
redistribute {rip | bgp | static | connected} [metric 0-16777214] [metric-type {1 | 2}] [tag 0-4294967295] [subnets]
no redistribute {rip | bgp | static | connected} [metric] [metric-type] [tag] [subnets]
```

Default

- metric—unspecified
- type—2
- tag—0

Command Mode

Router OSPF Config

5.2.1.40 **maximum-paths (OSPF)**

This command sets the number of paths that OSPF can report for a given destination where *maxpaths* is platform dependent.

Use the no command to reset the number of paths that OSPF can report for a given destination back to its default value.

```
maximum-paths maxpaths
no maximum-paths
```

Default

4

Command Mode

Router OSPF Config

5.2.1.41 **passive-interface default (OSPF)**

Use this command to enable global passive mode by default for all interfaces. It overrides any interface level passive mode. OSPF will not form adjacencies over a passive interface.

Use the no command to disable the global passive mode by default for all interfaces. Any interface previously configured to be passive reverts to nonpassive mode.

```
passive-interface default
no passive-interface default
```

Default

Disabled.

Command Mode

Router OSPF Config

5.2.1.42 **passive-interface (OSPF)**

Use this command to set the interface as passive. It overrides the global passive mode that is currently effective on the interface. The argument *unit/slot/port* corresponds to a physical routing interface or VLAN routing interface. The keyword *vlan* is used to specify the VLAN ID of the routing VLAN directly instead of a *unit/slot/port* format.

Use the no command to set the interface as nonpassive. It overrides the global passive mode that is currently effective on the interface.

```
passive-interface {unit/slot/port | vlan 1-4093}
no passive-interface {unit/slot/port | vlan 1-4093}
```

Default

Disabled.

Command Mode

Router OSPF Config

5.2.1.43 **timers pacing flood**

To adjust the rate at which OSPFv2 sends LS Update packets, use the `timers pacing flood` command in router OSPFv2 Global Configuration mode. OSPF distributes routing information in Link State Advertisements (LSAs), which are bundled into Link State Update (LS Update) packets. To reduce the likelihood of sending a neighbor more packets than it can buffer, OSPF rate limits the transmission of LS Update packets. By default, OSPF sends up to 30 updates per second on each interface (1/ the pacing interval). Use this command to adjust this packet rate.

Use the `no` command to revert LSA transmit pacing to the default rate.

```
timers pacing flood milliseconds
no timers pacing flood
```

Parameters

Parameter	Description
<i>milliseconds</i>	The average time between transmission of LS Update packets. The range is from 5 ms to 100 ms. The default is 33 ms.

Default

33 milliseconds.

Command Mode

OSPFv2 Router Config

5.2.1.44 **timers pacing lsa-group**

To adjust how OSPF groups LSAs for periodic refresh, use the `timers pacing lsa-group` command in OSPFv2 Router Configuration mode. OSPF refreshes self-originated LSAs approximately once every 30 minutes. When OSPF refreshes LSAs, it considers all self-originated LSAs whose age is from 1800 to 1800 plus the pacing group size. Grouping LSAs for refresh allows OSPF to combine refreshed LSAs into a minimal number of LS Update packets. Minimizing the number of Update packets makes LSA distribution more efficient.

When OSPF originates a new or changed LSA, it selects a random refresh delay for the LSA. When the refresh delay expires, OSPF refreshes the LSA. By selecting a random refresh delay, OSPF avoids refreshing a large number of LSAs at one time, even if a large number of LSAs are originated at one time.

```
timers pacing lsa-group seconds
```

Parameters

Parameter	Description
<i>seconds</i>	Width of the window in which LSAs are refreshed. The range for the pacing group window is from 10 to 1800 seconds.

Default

60 seconds.

Command Mode

OSPFv2 Router Config

5.2.1.45 timers spf

Use this command to configure the SPF delay time and hold time. The valid range for both parameters is 0-65535 seconds.

```
timers spf delay-time hold-time
```

Default

- delay-time - 5
- hold-time - 10

Command Mode

Router OSPF Config

5.2.1.46 Router OSPF Config

Use this command to enable individual OSPF traps, enable a group of trap flags at a time, or enable all the trap flags at a time. The different groups of trapflags, and each group's specific trapflags to enable or disable, are listed in below table.

- To enable the individual flag, enter the group name followed by that particular flag.
- To enable all the flags in that group, give the group name followed by `all`.
- To enable all the flags, give the command as `trapflags all`.

Use the no command to revert to the default reference bandwidth.

- To disable the individual flag, enter the group name followed by that particular flag.
- To disable all the flags in that group, give the group name followed by `all`.
- To disable all the flags, give the command as `trapflags all`.

Group	Flags
errors	<ul style="list-style-type: none">■ authentication-failure■ bad-packet■ config-error■ virt-authentication-failure■ virt-bad-packet■ virt-config-error
lsa	<ul style="list-style-type: none">■ lsa-maxage■ lsa-originate
overflow	<ul style="list-style-type: none">■ lsdbs-overflow■ lsdbs-approaching-overflow
retransmit	<ul style="list-style-type: none">■ packets■ virt-packets
state-change	<ul style="list-style-type: none">■ if-state-change■ neighbor-state-change■ virtif-state-change■ virtneighbor-state-change

```
trapflags {all | errors {all | authentication-failure |  
bad-packet | config-error | virtauthentication-failure |  
virt-bad-packet | virt-config-error} | lsa {all | lsa-max-  
age | lsa-originate} | overflow {all | lsdbs-overflow |  
lsdbs-approaching-overflow} | retransmit {all | packets |  
virt-packets} | state-change {all | if-state-change |  
neighbor-state-change | virtif-state-change | virtneighbor-  
state-change}}
```

```
no trapflags {all | errors {all | authentication-failure |
bad-packet | config-error | virtauthentication-failure |
virt-bad-packet | virt-config-error} | lsa {all | lsa-max-
age | lsa-originate} | overflow {all | lsdb-overflow |
lsdb-approaching-overflow} | retransmit {all | packets |
virt-packets} | state-change {all | if-state-change |
neighbor-state-change | virtif-statechange | virtneighbor-
state-change}}
```

Default

Disabled.

Command Mode

Router OSPF Config

5.2.2 OSPF Interface Commands

5.2.2.1 ip ospf area

Use this command to enable OSPFv2 and set the area ID of an interface or range of interfaces. The *area-id* is an IP address formatted as a 4-digit dotted-decimal number or a decimal value in the range of 0-4294967295. This command supersedes the effects of the *network area* command. It can also be used to configure the advertiseability of the secondary addresses on this interface into the OSPFv2 domain.

Use the no command to disable OSPF on an interface.

```
ip ospf area area-id [secondaries none]
no ip ospf area [secondaries none]
```

Default

Disabled.

Command Mode

Interface Config

5.2.2.2 bandwidth

By default, OSPF computes the link cost of an interface as the ratio of the reference bandwidth to the interface bandwidth. Reference bandwidth is specified with the auto-cost command. For the purpose of the OSPF link cost calculation, use the bandwidth command to specify the interface bandwidth. The bandwidth is specified in kilobits per second. If no bandwidth is configured, the bandwidth defaults to the actual interface bandwidth for port-based routing interfaces and to 10 Mbps for VLAN routing interfaces. This command does not affect the actual speed of an interface. You can use this command to configure a single interface or a range of interfaces.

Use the no command to set the interface bandwidth to its default value.

```
bandwidth 1-10000000
no bandwidth
```

Default

Actual interface bandwidth.

Command Mode

Interface Config

5.2.2.3 ip ospf authentication

This command sets the OSPF Authentication Type and Key for the specified interface or range of interfaces. The value of type is either none, simple or encrypt. The *key* is composed of standard displayable, noncontrol keystrokes from a Standard 101/102-key keyboard. The authentication key must be 8 bytes or less if the authentication type is simple. If the type is encrypt, the key may be up to 16 bytes. If the type is encrypt a *keyid* in the range of 0 and 255 must be specified. Unauthenticated interfaces do not need an authentication key or authentication key ID. There is no default value for this command.

Use the no command to set the default OSPF Authentication Type for the specified interface.

```
ip ospf authentication {none | {simple key} | {encrypt key  
keyid}}  
no ip ospf authentication
```

Command Mode

Interface Config

5.2.2.4 ip ospf cost

This command configures the cost on an OSPF interface or range of interfaces. The cost parameter has a range of 1 to 65535.

Use the no command to configure the default cost on an OSPF interface.

```
ip ospf cost 1-65535  
no ip ospf cost
```

Default

10

Command Mode

Interface Config

5.2.2.5 ip ospf database-filter all out

Use the command in Interface Configuration mode to disable OSPFv2 LSA flooding on an interface.

Use the no command in Interface Configuration mode to enable OSPFv2 LSA flooding on an interface.

```
ip ospf database-filter all out  
no ip ospf database-filter all out
```

Default

Disabled.

Command Mode

Interface Config

5.2.2.6 ip ospf dead-interval

This command sets the OSPF dead interval for the specified interface or range of interfaces. The value for seconds (range: 1-65535) is a valid positive integer, which represents the length of time in seconds that a router's Hello packets have not been seen before its neighbor routers declare that the router is down. The value for the length of time must be the same for all routers attached to a common network. This value should be some multiple of the Hello Interval (i.e. 4). Valid values range in seconds from 1 to 65535.

Use the no command to set the default OSPF dead interval for the specified interface.

```
ip ospf dead-interval seconds
no ip ospf dead-interval
```

Default

40

Command Mode

Interface Config

5.2.2.7 ip ospf hello-interval

This command sets the OSPF hello interval for the specified interface or range of interfaces. The value for seconds is a valid positive integer, which represents the length of time in seconds. The value for the length of time must be the same for all routers attached to a network. Valid values range from 1 to 65535.

Use the no command to sets the default OSPF hello interval for the specified interface.

```
ip ospf hello-interval seconds
no ip ospf hello-interval
```

Default

10

Command Mode

Interface Config

5.2.2.8 ip ospf network

Use this command to configure OSPF to treat an interface or range of interfaces as a point-to-point rather than broadcast interface. The `broadcast` option sets the OSPF network type to broadcast. The `point-to-point` option sets the OSPF network type to point-to-point. OSPF treats interfaces as broadcast interfaces by default. (Loopback interfaces have a special loopback network type, which cannot be changed.) When there are only two routers on the network, OSPF can operate more efficiently by treating the network as a point-to-point network. For point-to-point networks, OSPF does not elect a designated router or generate a network link state advertisement (LSA). Both endpoints of the link must be configured to operate in point-to-point mode.

Use the no command to return the OSPF network type to the default.

```
ip ospf network {broadcast | point-to-point}
no ip ospf network
```

Default

Broadcast.

Command Mode

Interface Config

5.2.2.9 ip ospf prefix-suppression

This command suppresses the advertisement of the IPv4 prefixes that are associated with an interface, except for those associated with secondary IPv4 addresses. This command takes precedence over the global configuration. If this configuration is not specified, the global prefix-suppression configuration applies.

prefix-suppression can be disabled at the interface level by using the disable option. The disable option is useful for excluding specific interfaces from performing prefix-suppression when the feature is enabled globally.

Note that the disable option disable is not equivalent to not configuring the interface specific prefix-suppression. If prefix-suppression is not configured at the interface level, the global prefix-suppression configuration is applicable for the IPv4 prefixes associated with the interface.

Use the no command to remove prefix-suppression configurations at the interface level. When `no ip ospf prefixsuppression` command is used, global prefix-suppression applies to the interface. Not configuring the command is not equal to disabling interface level prefix-suppression.

```
ip ospf prefix-suppression [disable]
no ip ospf prefix-suppression
```

Default

Prefix-suppression is not configured.

Command Mode

Interface Config

5.2.2.10 ip ospf priority

This command sets the OSPF priority for the specified router interface or range of interfaces. The priority of the interface is a priority integer from 0 to 255. A value of 0 indicates that the router is not eligible to become the designated router on this network.

Use the no command to set the default OSPF priority for the specified router interface.

```
ip ospf priority 0-255
no ip ospf priority
```

Default

1, which is the highest router priority.

Command Mode

Interface Config

5.2.2.11 ip ospf retransmit-interval

This command sets the OSPF retransmit Interval for the specified interface or range of interfaces. The retransmit interval is specified in seconds. The value for *seconds* is the number of seconds between link-state advertisement retransmissions for adjacencies belonging to this router interface. This value is also used when retransmitting database description and link-state request packets. Valid values range from 0 to 3600 (1 hour).

Use the no command to set the default OSPF retransmit Interval for the specified interface.

```
ip ospf retransmit-interval 0-3600
no ip ospf retransmit-interval
```

Default

5

Command Mode

Interface Config

5.2.2.12 ip ospf transmit-delay

This command sets the OSPF Transit Delay for the specified interface or range of interfaces. The transmit delay is specified in seconds. In addition, it sets the estimated number of seconds it takes to transmit a link state update packet over this interface. Valid values for *seconds* range from 1 to 3600 (1 hour).

Use the no command to set the default OSPF Transit Delay for the specified interface.

```
ip ospf transmit-delay 1-3600
no ip ospf transmit-delay
```

Default

1

Command Mode

Interface Config

5.2.2.13 ip ospf mtu-ignore

This command disables OSPF maximum transmission unit (MTU) mismatch detection on an interface or range of interfaces. OSPF Database Description packets specify the size of the largest IP packet that can be sent without fragmentation on the interface. When a router receives a Database Description packet, it examines the MTU advertised by the neighbor. By default, if the MTU is larger than the router can accept, the Database Description packet is rejected and the OSPF adjacency is not established.

Use the no command to enable the OSPF MTU mismatch detection.

```
ip ospf mtu-ignore
no ip ospf mtu-ignore
```

Default

Enabled.

Command Mode

Interface Config

5.2.3 IP Event Dampening Commands

5.2.3.1 dampening

Use this command to enable IP event dampening on a routing interface.

Use the no command to disable IP event dampening on a routing interface.

```
dampening [half-life period] [reuse-threshold suppress-  
threshold max-suppress-time [restart restart-penalty]]  
no dampening
```

Parameters

Parameter	Description
<code>half-life period</code>	The number of seconds it takes for the penalty to reduce by half. The configurable range is 1-30 seconds. Default value is 5 seconds.
<code>reuse-threshold</code>	The value of the penalty at which the dampened interface is restored. The configurable range is 1-20,000. Default value is 1000.
<code>suppress-threshold</code>	The value of the penalty at which the interface is dampened. The configurable range is 1-20,000. Default value is 2000.
<code>max-suppress-time</code>	The maximum amount of time (in seconds) an interface can be in suppressed state after it stops flapping. The configurable range is 1-255 seconds. The default value is four times of half-life period. If half-period value is allowed to default, the maximum suppress time defaults to 20 seconds.
<code>restart restart-penalty</code>	Penalty applied to the interface after the device reloads. The configurable range is 1-20,000. Default value is 2000.

Command Mode

Interface Config

5.2.3.2 show dampening interface

This command summarizes the number of interfaces configured with dampening and the number of interfaces being suppressed.

```
show dampening interface
```

Command Mode

Privileged EXEC

Example

The following shows example CLI display output for the command.

```
(Router)#show dampening interface
```

```
2 interfaces are configured with dampening.  
1 interface is being suppressed.
```

5.2.3.3 show interface dampening

This command displays the status and configured parameters of the interfaces configured with dampening.

```
show interface dampening
```

Command Mode

Privileged EXEC

Example

The following shows example CLI display output for the command.

```
(Router)# show interface dampening
```

```
Interface 0/2
```

```

Flaps Penalty Supp ReuseTm HalfL ReuseV
0 0 FALSE 0 5 1000
SuppV MaxSTm MaxP Restart
2000 20 16000 0

```

```

Interface 0/3

```

```

Flaps Penalty Supp ReuseTm HalfL ReuseV
6 1865 TRUE 18 20 1000
SuppV MaxSTm MaxP Restart
2001 30 2828 1500

```

Display Parameters

Term	Definition
Flaps	The number times the link state of an interface changed from UP to DOWN.
Penalty	Accumulated Penalty.
Supp	Indicates if the interface is suppressed or not.
ReuseTm	Number of seconds until the interface is allowed to come up again.
HalfL	Configured half-life period.
ReuseV	Configured reuse-threshold.
SuppV	Configured suppress threshold.
MaxSTm	Configured maximum suppress time in seconds.
MaxP	Maximum possible penalty.
Restart	Configured restart penalty.

- Note!**
1. The CLI command `clear counters` resets the flap count to zero.
 2. The interface CLI command `no shutdown` resets the suppressed state to `False`.
 3. Any change in the dampening configuration resets the current penalty, reuse time and suppressed state to their default values, meaning 0, 0, and `FALSE` respectively.



5.2.4 OSPF Graceful Restart Commands

The OSPF protocol can be configured to participate in the checkpointing service, so that these protocols can execute a “graceful restart” when the management unit fails. In a graceful restart, the hardware continues forwarding IPv4 packets using OSPF routes while a backup switch takes over management unit responsibility.

Graceful restart uses the concept of “helpful neighbors”. A fully adjacent router enters helper mode when it receives a link state announcement (LSA) from the restarting management unit indicating its intention of performing a graceful restart. In helper mode, a switch continues to advertise to the rest of the network that they have full adjacencies with the restarting router, thereby avoiding announcement of a topology change and the potential for flooding of LSAs and shortest-path-first (SPF) runs (which determine OSPF routes). Helpful neighbors continue to forward packets through the restarting router. The restarting router relearns the network topology from its helpful neighbors.

Graceful restart can be enabled for either planned or unplanned restarts, or both. A planned restart is initiated by the operator through the management command `initiate failover`. The operator may initiate a failover in order to take the management unit out of service (for example, to address a partial hardware failure), to correct faulty system behavior which cannot be corrected through less severe management

actions, or other reasons. An unplanned restart is an unexpected failover caused by a fatal hardware failure of the management unit or a software hang or crash on the management unit.

5.2.4.1 **nsf**

Use this command to enable the OSPF graceful restart functionality on an interface. To disable graceful restart, use the no form of the command.

Use the no command to disable graceful restart for all restarts.

```
nsf [ietf] [planned-only]
no nsf [ietf] [planned-only]
```

Parameters

Parameter	Description
<code>ietf</code>	(Optional) This keyword is accepted but not required.
<code>planned-only</code>	(Optional) This optional keyword indicates that OSPF should only perform a graceful restart when the restart is planned (i.e., when the restart is a result of the initiate failover command).

Default

Disabled.

Command Mode

OSPF Router Config

5.2.4.2 **nsf restart-interval**

Use this command to configure the number of seconds that the restarting router asks its neighbors to wait before exiting helper mode. This is referred to as the grace period. The restarting router includes the grace period in its grace LSAs. For planned restarts (using the initiate failover command), the grace LSAs are sent prior to restarting the management unit, whereas for unplanned restarts, they are sent after reboot begins.

The grace period must be set long enough to allow the restarting router to reestablish all of its adjacencies and complete a full database exchange with each of those neighbors.

Use the no command to revert the grace period to its default value.

```
nsf [ietf] restart-interval 1-1800
no nsf [ietf] restart-interval
```

Parameters

Parameter	Description
<code>ietf</code>	This keyword is accepted but not required.
<code>seconds</code>	The number of seconds that the restarting router asks its neighbors to wait before exiting helper mode. The range is from 1 to 1800 seconds.

Default

120 seconds.

Command Mode

OSPF Router Config

5.2.4.3 **nsf helper**

Use this command to enable helpful neighbor functionality for the OSPF protocol. You can enable this functionality for planned or unplanned restarts, or both.

Use the no command to disable helpful neighbor functionality for OSPF.

```
nsf helper [planned-only]
```

```
no nsf helper
```

Parameters

Parameter	Description
<code>planned-only</code>	This optional keyword indicates that OSPF should only help a restarting router performing a planned restart.

Default

OSPF may act as a helpful neighbor for both planned and unplanned restarts.

Command Mode

OSPF Router Config

5.2.4.4 **nsf ietf helper disable**

Use this command to disable helpful neighbor functionality for OSPF.

Note! *The commands `no nsf helper` and `nsf ietf helper disable` are functionally equivalent. The command `nsf ietf helper disable` is supported solely for compatibility with other network software CLI.*



```
nsf ietf helper disable
```

Command Mode

OSPF Router Config

5.2.4.5 **nsf helper strict-lsa-checking**

The restarting router is unable to react to topology changes. In particular, the restarting router will not immediately update its forwarding table; therefore, a topology change may introduce forwarding loops or black holes that persist until the graceful restart completes. By exiting the graceful restart on a topology change, a router tries to eliminate the loops or black holes as quickly as possible by routing around the restarting router. A helpful neighbor considers a link down with the restarting router to be a topology change, regardless of the strict LSA checking configuration.

Use this command to require that an OSPF helpful neighbor exit helper mode whenever a topology change occurs.

Use the no command to allow OSPF to continue as a helpful neighbor in spite of topology changes.

```
nsf [ietf] helper strict-lsa-checking
no nsf [ietf] helper strict-lsa-checking
```

Parameters

Parameter	Description
<code>ietf</code>	This keyword is accepted but not required.

Default

Enabled.

Command Mode

OSPF Router Config

5.2.5 OSPFv2 Stub Router Commands

5.2.5.1 max-metric router-lsa

To configure OSPF to enter stub router mode, use this command in Router OSPF Global Configuration mode. When OSPF is in stub router mode, as defined by RFC 3137, OSPF sets the metric in the nonstub links in its router LSA to LsInfinity. Other routers therefore compute very long paths through the stub router, and prefer any alternate path. Doing so eliminates all transit traffic through the stub router, when alternate routes are available. Stub router mode is useful when adding or removing a router from a network or to avoid transient routes when a router reloads.

You can administratively force OSPF into stub router mode. OSPF remains in stub router mode until you take OSPF out of stub router mode. Alternatively, you can configure OSPF to start in stub router mode for a configurable period of time after the router boots up.

If you set the summary LSA metric to 16,777,215, other routers will skip the summary LSA when they compute routes.

If you have configured the router to enter stub router mode on startup (max-metric router-lsa on-startup), and then enter max-metric router lsa, there is no change. If OSPF is administratively in stub router mode (the maxmetric router-lsa command has been given), and you configure OSPF to enter stub router mode on startup (max-metric router-lsa on-startup), OSPF exits stub router mode (assuming the startup period has expired) and the configuration is updated.

Use the no command to in OSPFv2 Router Configuration mode to disable stub router mode. The command clears either type of stub router mode (always or on-startup) and resets the summary-lsa option. If OSPF is configured to enter global configuration mode on startup, and during normal operation you want to immediately place OSPF in stub router mode, issue the command no max-metric router-lsa on-startup. The command no max-metric router-lsa summary-lsa causes OSPF to send summary LSAs with metrics computed using normal procedures defined in RFC 2328.

```
max-metric router-lsa [on-startup seconds] [summary-lsa  
{metric}]  
no max-metric router-lsa [on-startup] [summary-lsa]
```

Parameters

Parameter	Description
<i>on-startup</i>	(Optional) OSPF starts in stub router mode after a reboot.
<i>seconds</i>	(Required if on-startup) The number of seconds that OSPF remains in stub router mode after a reboot. The range is 5 to 86,400 seconds. There is no default value.
<i>summary-lsa</i>	(Optional) Set the metric in type 3 and type 4 summary LSAs to LsInfinity (0xFFFFFFFF).
<i>metric</i>	(Optional) Metric to send in summary LSAs when in stub router mode. The range is 1 to 16,777,215. The default is 16,711,680 (0xFF0000).

Default

OSPF is not in stub router mode by default.

Command Mode

OSPFv2 Router Config

5.2.5.2 clear ip ospf stub-router

Use the clear ip ospf stub-router command in Privileged EXEC mode to force OSPF to exit stub router mode for the specified virtual router when it has automatically entered stub router mode because of a resource limitation. OSPF only exits stub router mode if it entered stub router mode because of a resource limitation or if it is in stub router mode at startup. If no virtual router is specified, the command is executed for the default router. This command has no effect if OSPF is configured to be in stub router mode permanently.

```
clear ip ospf stub-router [vrf vrf-name]
```

Command Mode

Privileged EXEC

5.2.6 OSPF Show Commands

5.2.6.1 show ip ospf

This command displays OSPF global configuration information for the specified virtual router. If no router is specified, it displays information for the default router.

```
show ip ospf [vrf vrf-name]
```

Command Mode

Global Config

Example

The following shows example CLI display output for the command.

```
(alpha3) #show ip ospf
```

```
Router ID.....3.3.3.3
OSPF Admin Mode.....Enable
RFC 1583 Compatibility.....Enable
External LSDB Limit.....No Limit
Exit Overflow Interval.....0
Spf Delay Time.....5
Spf Hold Time.....10
Flood Pacing Interval.....33 ms
LSA Refresh Group Pacing Time.....60 sec
Opaque Capability.....Enable
AutoCost Ref BW.....100 Mbps
Default Passive Setting.....Disabled
Maximum Paths.....4
Default Metric.....Not configured
Stub Router Configuration.....<val>
Stub Router Startup Time.....<val> seconds
Summary LSA Metric Override.....Enabled (<met>)
Default Route Advertise.....Disabled
Always.....FALSE
Metric.....Not configured
Metric Type.....External Type 2
Number of Active Areas.....1 (1 normal, 0 stub, 0 nssa)
ABR Status.....Disable
ASBR Status.....Disable
Stub Router.....FALSE
Stub Router Status.....Inactive
Stub Router Reason.....<reason>
Stub Router Startup Time Remaining.....<duration> seconds
Stub Router Duration.....<duration>
External LSDB Overflow.....FALSE
```

```

External LSA Count.....0
External LSA Checksum.....0
AS_OPAQUE LSA Count.....0
AS_OPAQUE LSA Checksum.....0
New LSAs Originated.....55
LSAs Received.....82
LSA Count.....1
Maximum Number of LSAs.....24200
LSA High Water Mark.....9
AS Scope LSA Flood List Length.....0
Retransmit List Entries.....0
Maximum Number of Retransmit Entries....96800
Retransmit Entries High Water Mark.....1
NSF Helper Support.....Always
NSF Helper Strict LSA Checking.....Enabled
Prefix-suppression.....Disabled

```

Display Parameters

Note! *Some of the information below displays only if you enable OSPF and configure certain features.*



Term	Definition
Router ID	A 32-bit integer in dotted decimal format identifying the router, about which information is displayed. This is a configured value.
OSPF Admin Mode	Shows whether the administrative mode of OSPF in the router is enabled or disabled. This is a configured value.
RFC 1583 Compatibility	Indicates whether 1583 compatibility is enabled or disabled. This is a configured value.
External LSDB Limit	The maximum number of nondefault AS-external-LSA (link state advertisement) entries that can be stored in the link-state database.
Exit Overflow Interval	The number of seconds that, after entering overflow state, a router will attempt to leave overflow state.
Spf Delay Time	The number of seconds between two subsequent changes of LSAs, during which time the routing table calculation is delayed.
Spf Hold Time	The number of seconds between two consecutive spf calculations.
Flood Pacing Interval	The average time, in milliseconds, between LS Update packet transmissions on an interface. This is the value configured with the command “timers pacing flood”.
LSA Refresh Group Pacing Time	The size in seconds of the LSA refresh group window. This is the value configured with the command “timers pacing lsa-group” on page 932.
Opaque Capability	Shows whether the router is capable of sending Opaque LSAs. This is a configured value.
Autocost Ref BW	Shows the value of auto-cost reference bandwidth configured on the router.
Default Passive Setting	Shows whether the interfaces are passive by default.
Maximum Paths	The maximum number of paths that OSPF can report for a given destination.
Default Metric	Default value for redistributed routes.

Term	Definition
Stub Router Configuration	When OSPF runs out of resources to store the entire link state database, or any other state information, OSPF goes into stub router mode. As a stub router, OSPF reoriginates its own router LSAs, setting the cost of all nonstub interfaces to infinity. Use this field to set stub router configuration to one of Always , Startup , None .
Stub Router Startup Time	Configured value in seconds. This row is only listed if OSPF is configured to be a stub router at startup.
Summary LSA Metric Override	One of Enabled (<i>met</i>), Disabled , where <i>met</i> is the metric to be sent in summary LSAs when in stub router mode.
BFD Enabled	Displays the BFD status.
Default Route Advertise	Indicates whether the default routes received from other source protocols are advertised or not.
Always	Shows whether default routes are always advertised.
Metric	The metric of the routes being redistributed. If the metric is not configured, this field is blank.
Metric Type	Shows whether the routes are External Type 1 or External Type 2.
Number of Active Areas	The number of active OSPF areas. An "active" OSPF area is an area with at least one interface up.
ABR Status	Shows whether the router is an OSPF Area Border Router.
ASBR Status	Reflects whether the ASBR mode is enabled or disabled. Enable implies that the router is an autonomous system border router. The router automatically becomes an ASBR when it is configured to redistribute routes learnt from other protocols. The possible values for the ASBR status is enabled (if the router is configured to redistribute routes learned by other protocols) or disabled (if the router is not configured for the same).
Stub Router Status	One of Active , Inactive .
Stub Router Reason	One of Configured, Startup, Resource Limitation. Note: The row is only listed if stub router is active.
Stub Router Startup Time Remaining	The remaining time, in seconds, until OSPF exits stub router mode. This row is only listed if OSPF is in startup stub router mode.
Stub Router Duration	The time elapsed since the router last entered the stub router mode. The row is only listed if stub router is active and the router entered stub mode because of a resource limitation. The duration is displayed in DD:HH:MM:SS format.
External LSDB Overflow	When the number of nondefault external LSAs exceeds the configured limit, External LSDB Limit, OSPF goes into LSDB overflow state. In this state, OSPF withdraws all of its selforiginated nondefault external LSAs. After the Exit Overflow Interval, OSPF leaves the overflow state, if the number of external LSAs has been reduced.
External LSA Count	The number of external (LS type 5) link-state advertisements in the link-state database.
External LSA Checksum	The sum of the LS checksums of external link-state advertisements contained in the linkstate database.
AS_OPAQUE LSA Count	Shows the number of AS Opaque LSAs in the link-state database.
AS_OPAQUE LSA Checksum	Shows the sum of the LS Checksums of AS Opaque LSAs contained in the link-state database.
New LSAs Originated	The number of new link-state advertisements that have been originated.
LSAs Received	The number of link-state advertisements received determined to be new instantiations.

Term	Definition
LSA Count	The total number of link state advertisements currently in the link state database.
Maximum Number of LSAs	The maximum number of LSAs that OSPF can store.
LSA High Water Mark	The maximum size of the link state database since the system started.
AS Scope LSA Flood List Length	The number of LSAs currently in the global flood queue waiting to be flooded through the OSPF domain. LSAs with AS flooding scope, such as type 5 external LSAs and type 11 Opaque LSAs.
Retransmit List Entries	The total number of LSAs waiting to be acknowledged by all neighbors. An LSA may be pending acknowledgment from more than one neighbor.
Maximum Number of Retransmit Entries	The maximum number of LSAs that can be waiting for acknowledgment at any given time.
Retransmit Entries High Water Mark	The maximum number of LSAs on all neighbors' retransmit lists at any given time.
NSF Support	Indicates whether nonstop forwarding (NSF) is enabled for the OSPF protocol for planned restarts, unplanned restarts or both ("Always").
NSF Restart Interval	The user-configurable grace period during which a neighboring router will be in the helper state after receiving notice that the management unit is performing a graceful restart.
NSF Restart Status	The current graceful restart status of the router. <ul style="list-style-type: none"> ■ Not Restarting ■ Planned Restart ■ Unplanned Restart
NSF Restart Age	Number of seconds until the graceful restart grace period expires.
NSF Restart Exit Reason	Indicates why the router last exited the last restart: <ul style="list-style-type: none"> ■ None - Graceful restart has not been attempted. ■ In Progress - Restart is in progress. ■ Completed - The previous graceful restart completed successfully. ■ Timed Out - The previous graceful restart timed out. ■ Topology Changed - The previous graceful restart terminated prematurely because of a topology change.
NSF Help Support	Indicates whether helpful neighbor functionality has been enabled for OSPF for planned restarts, unplanned restarts, or both (Always).
NSF help Strict LSA checking	Indicates whether strict LSA checking has been enabled. If enabled, then an OSPF helpful neighbor will exit helper mode whenever a topology change occurs. If disabled, an OSPF neighbor will continue as a helpful neighbor in spite of topology changes.
Prefix-suppression	Displays whether prefix-suppression is enabled or disabled.

5.2.6.2 show ip ospf abr

This command displays the internal OSPF routing table entries to Area Border Routers (ABR) for the specified virtual router. If no router is specified, it displays information for the default router.

```
show ip ospf abr [vrf vrf-name]
```

Command Mode

- Privileged EXEC
- User EXEC

Display Parameters

Term	Definition
Type	The type of the route to the destination. It can be either: <ul style="list-style-type: none">■ intra — Intra-area route■ inter — Inter-area route
Router ID	Router ID of the destination.
Cost	Cost of using this route.
Area ID	The area ID of the area from which this route is learned.
Next Hop	Next hop toward the destination.
Next Hop Intf	The outgoing router interface to use when forwarding traffic to the next hop.

5.2.6.3 show ip ospf area

This command displays information about the area for the specified virtual router. If no router is specified, it displays information for the default router. The *areaid* identifies the OSPF area that is being displayed.

```
show ip ospf area areaid [vrf vrf-name]
```

Command Mode

- Privileged EXEC
- User EXEC

Example

The following shows example CLI display output for the command.

```
(R1)#show ip ospf area 1

AreaID.....0.0.0.1
External Routing.....Import External LSAs
Spf Runs.....10
Area Border Router Count.....0
Area LSA Count.....3004
Area LSA Checksum.....0x5e0abed
Flood List Length.....0
Import Summary LSAs.....Enable
```

Display Parameters

Term	Definition
AreaID	The area id of the requested OSPF area.
External Routing	A number representing the external routing capabilities for this area.
Spf Runs	The number of times that the intra-area route table has been calculated using this area's link-state database.
Area Border Router Count	The total number of area border routers reachable within this area.

Term	Definition
Area LSA Count	Total number of link-state advertisements in this area's link-state database, excluding AS External LSA's.
Area LSA Checksum	A number representing the Area LSA Checksum for the specified AreaID excluding the external (LS type 5) link-state advertisements.
Flood List Length	The number of LSAs waiting to be flooded within the area.
Import Summary LSAs	Shows whether to import summary LSAs.
OSPF Stub Metric Value	The metric value of the stub area. This field displays only if the area is a configured as a stub area.

The following OSPF NSSA specific information displays only if the area is configured as an NSSA:

Parameter	Description
Import Summary LSAs	Shows whether to import summary LSAs into the NSSA.
Redistribute into NSSA	Shows whether to redistribute information into the NSSA.
Default Information Originate	Shows whether to advertise a default route into the NSSA.
Default Metric	Shows whether to advertise a default route into the NSSA.
Default Metric Type	The metric type for the default route advertised into the NSSA.
Translator Role	The NSSA translator role of the ABR, which is always or candidate.
Translator Stability Interval	The amount of time that an elected translator continues to perform its duties after it determines that its translator status has been deposed by another router.
Translator State	Shows whether the ABR translator state is disabled, always, or elected.

5.2.6.4 show ip ospf asbr

This command displays the internal OSPF routing table entries to Autonomous System Boundary Routers (ASBR) for the specified virtual router. If no router is specified, it displays information for the default router.

```
show ip ospf asbr [vrf vrf-name]
```

Command Mode

- Privileged EXEC
- User EXEC

Display Parameters

Term	Definition
Type	The type of the route to the destination. It can be one of the following values: <ul style="list-style-type: none"> ■ intra — Intra-area route ■ inter — Inter-area route
Router ID	Router ID of the destination.
Cost	Cost of using this route.
Area ID	The area ID of the area from which this route is learned.
Next Hop	Next hop toward the destination.
Next Hop Intf	The outgoing router interface to use when forwarding traffic to the next hop.

5.2.6.5 show ip ospf database

This command displays information about the link state database when OSPF is enabled for the specified virtual router. If no router is specified, it displays information for the default router. If you do not enter any parameters, the command displays the LSA headers for all areas. Use the optional *areaid* parameter to display database information about a specific area. Use the optional parameters to specify the type of link state advertisements to display.

```
show ip ospf [areaid] database [vrf vrf-name] [{database-  
summary | [{asbr-summary | external | network | nssa-exter-  
nal | opaque-area | opaque-as | opaque-link | router | sum-  
mary}]} [lsid] [{adv-router [ipaddr] | self-originate}}]}
```

Parameters

Parameter	Description
<i>vrf vrf-name</i>	(Optional) Specifies the virtual router for which to display information.
<i>asbr-summary</i>	(Optional) Use <i>asbr-summary</i> to show the autonomous system boundary router (ASBR) summary LSAs.
<i>external</i>	(Optional) Use <i>external</i> to display the external LSAs.
<i>network</i>	(Optional) Use <i>network</i> to display the network LSAs.
<i>nssa-external</i>	(Optional) Use <i>nssa-external</i> to display NSSA external LSAs.
<i>opaque-area</i>	(Optional) Use <i>opaque-area</i> to display area opaque LSAs.
<i>opaque-as</i>	(Optional) Use <i>opaque-as</i> to display AS opaque LSAs.
<i>opaque-link</i>	(Optional) Use <i>opaque-link</i> to display link opaque LSAs.
<i>router</i>	(Optional) Use <i>router</i> to display router LSAs.
<i>summary</i>	(Optional) Use <i>summary</i> to show the LSA database summary information.
<i>lsid</i>	(Optional) Use <i>lsid</i> to specify the link state ID (LSID). The value of <i>lsid</i> can be an IP address or an integer in the range of 0-4294967295.
<i>adv-router</i>	(Optional) Use <i>adv-router</i> to show the LSAs that are restricted by the advertising router.
<i>self-originate</i>	(Optional) Use <i>self-originate</i> to display the LSAs in that are self originated. The information below is only displayed if OSPF is enabled.

Command Mode

- Privileged EXEC
- User EXEC

Display Parameters

For each link-type and area, the following information is displayed:

Term	Definition
Link Id	A number that uniquely identifies an LSA that a router originates from all other self originated LSAs of the same LS type.
Adv Router	The Advertising Router. Is a 32-bit dotted decimal number representing the LSDB interface.
Age	A number representing the age of the link state advertisement in seconds.
Sequence	A number that represents which LSA is more recent.
Checksum	The total number LSA checksum.

Term	Definition
Options	This is an integer. It indicates that the LSA receives special handling during routing calculations.
Rtr Opt	Router Options are valid for router links only.

5.2.6.6 **show ip ospf database database-summary**

Use this command to display the number of each type of LSA in the database for each area and for the router. The command also displays the total number of LSAs in the database.

```
show ip ospf database database-summary
```

Command Mode

- Privileged EXEC
- User EXEC

Display Parameters

Term	Definition
Router	Total number of router LSAs in the OSPF link state database.
Network	Total number of network LSAs in the OSPF link state database.
Summary Net	Total number of summary network LSAs in the database.
Summary ASBR	Number of summary ASBR LSAs in the database.
Type-7 Ext	Total number of Type-7 external LSAs in the database.
Self-Originated Type-7	Total number of self originated AS external LSAs in the OSPF link state database.
Opaque Link	Number of opaque link LSAs in the database.
Opaque Area	Number of opaque area LSAs in the database.
Subtotal	Number of entries for the identified area.
Opaque AS	Number of opaque AS LSAs in the database.
Total	Number of entries for all areas.

5.2.6.7 **show ip ospf interface**

This command displays the information for the IFO object or virtual interface tables. The argument *unit/slot/port* corresponds to a physical routing interface or VLAN routing interface. The keyword *vlan* is used to specify the VLAN ID of the routing VLAN directly instead of a *unit/slot/port* format.

```
show ip ospf interface {unit/slot/port | vlan 1-4093 | loopback loopback-id}
```

Command Mode

- Privileged EXEC
- User EXEC

Example

The following shows example CLI display output for the command when the OSPF Admin Mode is disabled.

```
(Routing)#show ip ospf interface 1/0/1

IP Address.....0.0.0.0
Subnet Mask.....0.0.0.0
Secondary IP Address(es).....
OSPF Admin Mode.....Disable
OSPF Area ID.....0.0.0.0
OSPF Network Type.....Broadcast
```

```

Router Priority.....1
Retransmit Interval.....5
Hello Interval.....10
Dead Interval.....40
LSA Ack Interval.....1
Transmit Delay.....1
Authentication Type.....None
Metric Cost.....1 (computed)
Passive Status.....Non-passive interface
OSPF Mtu-ignore.....Disable
Flood Blocking.....Disable

```

```

OSPF is not enabled on this interface.
(Routing)#

```

Display Parameters

Term	Definition
IP Address	The IP address for the specified interface.
Subnet Mask	A mask of the network and host portion of the IP address for the OSPF interface.
Secondary IP Address(es)	The secondary IP addresses if any are configured on the interface.
OSPF Admin Mode	States whether OSPF is enabled or disabled on a router interface.
OSPF Area ID	The OSPF Area ID for the specified interface.
OSPF Network Type	The type of network on this interface that the OSPF is running on.
Router Priority	A number representing the OSPF Priority for the specified interface.
Retransmit Interval	A number representing the OSPF Retransmit Interval for the specified interface.
Hello Interval	A number representing the OSPF Hello Interval for the specified interface.
Dead Interval	A number representing the OSPF Dead Interval for the specified interface.
LSA Ack Interval	A number representing the OSPF LSA Acknowledgment Interval for the specified interface.
Transmit Delay	A number representing the OSPF Transmit Delay Interval for the specified interface.
Authentication Type	The OSPF Authentication Type for the specified interface are: none, simple, and encrypt.
Metric Cost	The cost of the OSPF interface.
Passive Status	Shows whether the interface is passive or not.
OSPF MTU-ignore	Indicates whether to ignore MTU mismatches in database descriptor packets sent from neighboring routers.
Flood Blocking	Indicates whether flood blocking is enabled on the interface.

The information below will only be displayed if OSPF is enabled.

Term	Definition
OSPF Interface Type	Broadcast LANs, such as Ethernet and IEEE 802.5, take the value broadcast. The OSPF Interface Type will be 'broadcast'.
State	The OSPF Interface States are: down, loopback, waiting, point-to-point, designated router, and backup designated router.

Term	Definition
Designated Router	The router ID representing the designated router.
Backup Designated Router	The router ID representing the backup designated router.
Number of Link Events	The number of link events.
Local Link LSAs	The number of Link Local Opaque LSAs in the link-state database.
Local Link LSA Checksum	The sum of LS Checksums of Link Local Opaque LSAs in the link-state database.
Prefix-suppression	Displays whether prefix-suppression is enabled, disabled, or unconfigured on the given interface.

5.2.6.8 **show ip ospf interface brief**

This command displays brief information for the IFO object or virtual interface tables for the specified virtual router. If no router is specified, it displays information for the default router.

```
show ip ospf interface brief [vrf vrf-name]
```

Command Mode

- Privileged EXEC
- User EXEC

Display Parameters

Term	Definition
Interface	<i>unit/slot/port</i>
OSPF Admin Mode	States whether OSPF is enabled or disabled on a router interface.
OSPF Area ID	The OSPF Area Id for the specified interface.
Router Priority	A number representing the OSPF Priority for the specified interface.
Cost	The metric cost of the OSPF interface.
Hello Interval	A number representing the OSPF Hello Interval for the specified interface.
Dead Interval	A number representing the OSPF Dead Interval for the specified interface.
Retransmit Interval	A number representing the OSPF Retransmit Interval for the specified interface.
Interface Transmit Delay	A number representing the OSPF Transmit Delay for the specified interface.
LSA Ack Interval	A number representing the OSPF LSA Acknowledgment Interval for the specified interface.

5.2.6.9 show ip ospf interface stats

This command displays the statistics for a specific interface. The information below will only be displayed if OSPF is enabled. The argument `unit/slot/port` corresponds to a physical routing interface or VLAN routing interface. The keyword `vlan` is used to specify the VLAN ID of the routing VLAN directly instead of a `unit/slot/port` format.

```
show ip ospf interface stats {unit/slot/port | vlan 1-4093}
```

Command Mode

- Privileged EXEC
- User EXEC

Display Parameters

Term	Definition
OSPF Area ID	The area id of this OSPF interface.
Area Border Router Count	The total number of area border routers reachable within this area. This is initially zero, and is calculated in each SPF pass.
AS Border Router Count	The total number of Autonomous System border routers reachable within this area.
Area LSA Count	The total number of link-state advertisements in this area's link-state database, excluding AS External LSAs.
IP Address	The IP address associated with this OSPF interface.
OSPF Interface Events	The number of times the specified OSPF interface has changed its state, or an error has occurred.
Virtual Events	The number of state changes or errors that occurred on this virtual link.
Neighbor Events	The number of times this neighbor relationship has changed state, or an error has occurred.
Sent Packets	The number of OSPF packets transmitted on the interface.
Received Packets	The number of valid OSPF packets received on the interface.
Discards	The number of received OSPF packets discarded because of an error in the packet or an error in processing the packet.
Bad Version	The number of received OSPF packets whose version field in the OSPF header does not match the version of the OSPF process handling the packet.
Source Not On Local Subnet	The number of received packets discarded because the source IP address is not within a subnet configured on a local interface. Note: This field applies only to OSPFv2.
Virtual Link Not Found	The number of received OSPF packets discarded where the ingress interface is in a nonbackbone area and the OSPF header identifies the packet as belonging to the backbone, but OSPF does not have a virtual link to the packet's sender.
Area Mismatch	The number of OSPF packets discarded because the area ID in the OSPF header is not the area ID configured on the ingress interface.
Invalid Destination Address	The number of OSPF packets discarded because the packet's destination IP address is not the address of the ingress interface and is not the AllDrRouters or AllSpfRouters multicast addresses.
Wrong Authentication Type	The number of packets discarded because the authentication type specified in the OSPF header does not match the authentication type configured on the ingress interface. Note: This field applies only to OSPFv2.

Term	Definition
Authentication Failure	The number of OSPF packets dropped because the sender is not an existing neighbor or the sender's IP address does not match the previously recorded IP address for that neighbor. Note: This field applies only to OSPFv2.
No Neighbor at Source Address	The number of OSPF packets dropped because the sender is not an existing neighbor or the sender's IP address does not match the previously recorded IP address for that neighbor. Note: Does not apply to Hellos.
Invalid OSPF Packet Type	The number of OSPF packets discarded because the packet type field in the OSPF header is not a known type.
Hellos Ignored	The number of received Hello packets that were ignored by this router from the new neighbors after the limit has been reached for the number of neighbors on an interface or on the system as a whole.

The below table lists the number of OSPF packets of each type sent and received on the interface.

Packet Type	Sent	Received
Hello	6960	6960
Database Description	3	3
LS Request	1	1
LS Update	141	42
LS Acknowledgment	40	135

5.2.6.10 show ip ospf lsa-group

This command displays the number of self-originated LSAs within each LSA group for the specified virtual router. If no router is specified, it displays information for the default router.

```
show ip ospf lsa-group [vrf vrf-name]
```

Command Mode

- Privileged EXEC
- User EXEC

Display Parameters

Term	Definition
Total self-originated LSAs	The number of LSAs the router is currently originating.
Average LSAs per group	The number of self-originated LSAs divided by the number of LSA groups. The number of LSA groups is the refresh interval (1800 seconds) divided by the pacing interval (configured with timers pacing lsa-group) plus two.
Pacing group limit	The maximum number of self-originated LSAs in one LSA group. If the number of LSAs in a group exceeds this limit, OSPF redistributes LSAs throughout the refresh interval to achieve better balance.
Groups	For each LSA pacing group, the output shows the range of LSA ages in the group and the number of LSAs in the group.

5.2.6.11 show ip ospf neighbor

This command displays information about OSPF neighbors for the specified virtual router. If no router is specified, it displays information for the default router. If you do not specify a neighbor IP address, the output displays summary information in a table. If you specify an interface or tunnel, only the information for that interface or tunnel displays, if the interface is a physical routing interface and vlan format if the interface is a routing vlan. The *ipaddress* is the IP address of the neighbor, and when you specify this, detailed information about the neighbor displays. The information below only displays if OSPF is enabled and the interface has a neighbor.

```
show ip ospf neighbor [vrf vrf-name] [interface {unit/slot/
port | vlan 1-4093}] [ipaddress]
```

Command Mode

- Privileged EXEC
- User EXEC

Example

The following shows example CLI display output for the command.

```
(alpha1)#show ip ospf neighbor 170.1.1.50

Interface.....0/17
Neighbor IP Address.....170.1.1.50
Interface Index.....17
Area Id.....0.0.0.2
Options.....0x2
Router Priority.....1
Dead timer due in (secs).....15
Up Time.....0 days 2 hrs 8 mins 46 secs
State.....Full/BACKUP-DR
Events.....4
Retransmitted LSAs.....32
Retransmission Queue Length.....0
Restart Helper Status.....Helping
Restart Reason.....Software Restart (1)
Remaining Grace Time.....10 sec
Restart Helper Exit Reason.....In Progress
```

Display Parameters

If you do not specify an IP address, a table with the following columns displays for all neighbors or the neighbor associated with the interface that you specify:

Term	Definition
Router ID	The 4-digit dotted-decimal number of the neighbor router.
Priority	The OSPF priority for the specified interface. The priority of an interface is a priority integer from 0 to 255. A value of '0' indicates that the router is not eligible to become the designated router on this network.
IP Address	The IP address of the neighbor.
Interface	The interface of the local router in <i>unit/slot/port</i> format.

Term	Definition
State	<p>The state of the neighboring routers. Possible values are:</p> <ul style="list-style-type: none"> ■ Down - Initial state of the neighbor conversation; no recent information has been received from the neighbor. ■ Attempt - No recent information has been received from the neighbor but a more concerted effort should be made to contact the neighbor. ■ Init - An Hello packet has recently been seen from the neighbor, but bidirectional communication has not yet been established. ■ 2 way—Communication between the two routers is bidirectional. ■ Exchange start - The first step in creating an adjacency between the two neighboring routers, the goal is to decide which router is the master and to decide upon the initial DD sequence number. ■ Exchange - The router is describing its entire link state database by sending Database Description packets to the neighbor. ■ Loading - Link State Request packets are sent to the neighbor asking for the more recent LSAs that have been discovered (but not yet received) in the Exchange state. ■ Full - The neighboring routers are fully adjacent and they will now appear in router-LSAs and network-LSAs.
Dead Time	The amount of time, in seconds, to wait before the router assumes the neighbor is unreachable.

If you specify an IP address for the neighbor router, the following fields display:

Term	Definition
Interface	unit/slot/port
Neighbor IP Address	The IP address of the neighbor router.
Interface Index	The interface ID of the neighbor router.
Area ID	The area ID of the OSPF area associated with the interface.
Options	An integer value that indicates the optional OSPF capabilities supported by the neighbor. The neighbor's optional OSPF capabilities are also listed in its Hello packets. This enables received Hello Packets to be rejected (i.e., neighbor relationships will not even start to form) if there is a mismatch in certain crucial OSPF capabilities.
Router Priority	The OSPF priority for the specified interface. The priority of an interface is a priority integer from 0 to 255. A value of '0' indicates that the router is not eligible to become the designated router on this network.
Dead Timer Due	The amount of time, in seconds, to wait before the router assumes the neighbor is unreachable.
Up Time	Neighbor uptime; how long since the adjacency last reached the Full state.
State	The state of the neighboring routers.
Events	The number of times this neighbor relationship has changed state, or an error has occurred.
Retransmitted LSAs	The number of LSAs retransmitted to this neighbor.
Retransmission Queue Length	An integer representing the current length of the retransmission queue of the specified neighbor router Id of the specified interface.

Term	Definition
Restart Helper Status	<p data-bbox="606 208 1441 275">Indicates the status of this router as a helper during a graceful restart of the router specified in the command line:</p> <ul style="list-style-type: none"> <li data-bbox="606 275 1441 499">■ Helping - This router is acting as a helpful neighbor to this neighbor. A helpful neighbor does not report an adjacency change during graceful restart, but continues to advertise the restarting router as a FULL adjacency. A helpful neighbor continues to forward data packets to the restarting router, trusting that the restarting router's forwarding table is maintained during the restart. <li data-bbox="606 499 1441 533">■ Not Helping - This router is not a helpful neighbor at this time.
Restart Reason	<p data-bbox="606 544 1441 611">When this router is in helpful neighbor mode, this indicates the reason for the restart as provided by the restarting router:</p> <ul style="list-style-type: none"> <li data-bbox="606 611 1441 645">■ Unknown (0) <li data-bbox="606 645 1441 678">■ Software restart (1) <li data-bbox="606 678 1441 712">■ Software reload/upgrade (2) <li data-bbox="606 712 1441 745">■ Switch to redundant control processor (3) <li data-bbox="606 745 1441 779">■ Unrecognized - a value not defined in RFC 3623 <p data-bbox="606 790 1441 891">When FASTPATH sends a grace LSA, it sets the Restart Reason to Software Restart on a planned warm restart (when the initiate failover command is invoked), and to Unknown on an unplanned warm restart.</p>
Remaining Grace Time	<p data-bbox="606 902 1441 992">The number of seconds remaining the in current graceful restart interval. This is displayed only when this router is currently acting as a helpful neighbor for the router specified in the command.</p>
Restart Helper Exit Reason	<p data-bbox="606 1003 1441 1070">Indicates the reason that the specified router last exited a graceful restart.</p> <ul style="list-style-type: none"> <li data-bbox="606 1070 1441 1104">■ None - Graceful restart has not been attempted <li data-bbox="606 1104 1441 1137">■ In Progress - Restart is in progress <li data-bbox="606 1137 1441 1205">■ Completed - The previous graceful restart completed successfully <li data-bbox="606 1205 1441 1238">■ Timed Out - The previous graceful restart timed out <li data-bbox="606 1238 1441 1317">■ Topology Changed - The previous graceful restart terminated prematurely because of a topology change

5.2.6.12 show ip ospf range

This command displays the set of OSPFv2 area ranges configured for a given area for the specified virtual router. If no router is specified, it displays information for the default router.

```
show ip ospf range areaid [vrf vrf-name]
```

Command Mode

Privileged EXEC

Example

The following shows example CLI display output for the command.

```
(R1)#show ip ospf range 0
```

Prefix	Subnet Mask	Type	Action	Cost	Active
10.1.0.0	255.255.0.0	S	Advertise	Auto	N
172.20.0.0	255.255.0.0	S	Advertise	500	Y

Display Parameters

Term	Definition
Prefix	The summary prefix.
Subnet Mask	The subnetwork mask of the summary prefix.
Type	S (Summary Link) or E (External Link)
Action	Advertise or Suppress
Cost	Metric to be advertised when the range is active. If a static cost is not configured, the field displays Auto . If the action is Suppress , the field displays N/A .
Active	Whether the range is currently active. Y or N .

5.2.6.13 show ip ospf statistics

This command displays information about recent Shortest Path First (SPF) calculations for the specified virtual router. If no router is specified, it displays information for the default router. The SPF is the OSPF routing table calculation. The output lists the number of times the SPF has run for each OSPF area. A table follows this information. For each of the 15 most recent SPF runs, the command shows statistics for how long ago the SPF ran, how long the SPF took, the reasons why the SPF was scheduled, the individual components of the routing table calculation time and to show the RIB update time. The most recent statistics are displayed at the end of the table.

```
show ip ospf statistics [vrf vrf-name]
```

Command Mode

Privileged EXEC

Example

The following shows example CLI display output for the command.

```
(Router)#show ip ospf statistics
```

```
Area 0.0.0.0: SPF algorithm executed 15 times
Delta T   Intra  Summ   Ext   SPF Total   RIB Update   Reason
00:05:33  0      0      0     0           0             R
00:05:30  0      0      0     0           0             R
00:05:19  0      0      0     0           0             N, SN
00:05:15  0      10     0    10           0             R, N, SN
00:05:11  0      0      0     0           0             R
00:04:50  0      60     0    60          460           R, N
```

00:04:46	0	90	0	100	60	R, N
00:03:42	0	70	10	90	160	R
00:03:39	0	70	40	120	240	X
00:03:36	0	60	60	130	160	X
00:01:28	0	60	50	130	240	X
00:01:25	0	30	50	110	310	SN
00:01:22	0	0	40	50	260	SN
00:01:19	0	0	20	20	190	X
00:01:16	0	0	0	0	110	R, X

Display Parameters

Term	Definition
Delta T	The time since the routing table was computed. The time is in the format hours, minutes, and seconds (hh:mm:ss).
Intra	The time taken to compute intra-area routes, in milliseconds.
Summ	The time taken to compute inter-area routes, in milliseconds.
Ext	The time taken to compute external routes, in milliseconds.
SPF Total	The total time to compute routes, in milliseconds. The total may exceed the sum of the Intra, Summ, and Ext times.
RIB Update	The time from the completion of the routing table calculation until all changes have been made in the common routing table [the Routing Information Base (RIB)], in milliseconds.
Reason	The event or events that triggered the SPF. Reason codes are as follows: <ul style="list-style-type: none"> ■ R - new router LSA ■ N - new network LSA ■ SN - new network summary LSA ■ SA - new ASBR summary LSA ■ X - new external LSA

5.2.6.14 show ip ospf stub table

This command displays the OSPF stub table for the virtual router. If no router is specified, the information for the default router will be displayed. The information below will only be displayed if OSPF is initialized on the switch.

```
show ip ospf stub table [vrf vrf-name]
```

Command Mode

- Privileged EXEC
- User EXEC

Display Parameters

Term	Definition
Area ID	A 32-bit identifier for the created stub area.
Type of Service	The type of service associated with the stub metric. FASTPATH only supports Normal TOS.
Metric Val	The metric value is applied based on the TOS. It defaults to the least metric of the type of service among the interfaces to other areas. The OSPF cost for a route is a function of the metric value.
Import Summary LSA	Controls the import of summary LSAs into stub areas.

5.2.6.15 show ip ospf traffic

This command displays OSPFv2 packet and LSA statistics and OSPFv2 message queue statistics for the virtual router. If no router is specified, the information for the default router will be displayed. Packet statistics count packets and LSAs since OSPFv2 counters were last cleared (using the command “clear ip ospf counters”).

Note! The “clear ip ospf counters” command does not clear the message queue high water marks.



```
show ip ospf traffic [vrf vrf-name]
```

Command Mode

Privileged EXEC

Example

The following shows example CLI display output for the command.

```
(Router)#show ip ospf traffic
```

```
Time Since Counters Cleared: 4000 seconds
```

OSPFv2 Packet Statistics

	Hello	Database Desc	LS Request	LS Update	LS ACK	Total
Recd:	500	10	20	50	20	600
Sent:	400	8	16	40	16	480

```
LSAs Retransmitted.....0
```

```
LS Update Max Receive Rate.....20 pps
```

```
LS Update Max Send Rate.....10 pps
```

Number of LSAs Received

```
T1 (Router).....10
T2 (Network).....0
T3 (Net Summary).....300
T4 (ASBR Summary).....15
T5 (External).....20
T7 (NSSA External).....0
T9 (Link Opaque).....0
T10 (Area Opaque).....0
T11 (AS Opaque).....0
Total.....345
```

OSPFv2 Queue Statistics

	Current	Max	Drops	Limit
Hello	0	10	0	500
ACK	2	12	0	1680
Data	24	47	0	500
Event	1	8	0	1000

Display Parameters

Term	Definition
OSPFv2 Packet Statistics	The number of packets of each type sent and received since OSPF counters were last cleared.
LSAs Retransmitted	The number of LSAs retransmitted by this router since OSPF counters were last cleared.
LS Update Max Receive Rate	The maximum rate of LS Update packets received during any 5-second interval since OSPF counters were last cleared. The rate is in packets per second.
LS Update Max Send Rate	The maximum rate of LS Update packets transmitted during any 5-second interval since OSPF counters were last cleared. The rate is in packets per second.
Number of LSAs Received	The number of LSAs of each type received since OSPF counters were last cleared.
OSPFv2 Queue Statistics	For each OSPFv2 message queue, the current count, the high water mark, the number of packets that failed to be enqueued, and the queue limit. The high water marks are not cleared when OSPF counters are cleared.

5.2.6.16 `show ip ospf virtual-link`

This command displays the OSPF Virtual Interface information for a specific area and neighbor for the virtual router. If no router is specified, the information for the default router will be displayed. The `areaid` parameter identifies the area and the `neighbor` parameter identifies the neighbor's Router ID.

```
show ip ospf virtual-link [vrf vrf-name] areaid neighbor
```

Command Mode

- Privileged EXEC
- User EXEC

Display Parameters

Term	Definition
Area ID	The area id of the requested OSPF area.
Neighbor Router ID	The input neighbor Router ID.
Hello Interval	The configured hello interval for the OSPF virtual interface.
Dead Interval	The configured dead interval for the OSPF virtual interface.
Interface Transmit Delay	The configured transmit delay for the OSPF virtual interface.
Retransmit Interval	The configured retransmit interval for the OSPF virtual interface.
Authentication Type	The configured authentication type of the OSPF virtual interface.
State	The OSPF Interface States are: down, loopback, waiting, point-to-point, designated router, and backup designated router. This is the state of the OSPF interface.
Neighbor State	The neighbor state.

5.2.6.17 **show ip ospf virtual-link brief**

This command displays the OSPF Virtual Interface information for all areas in the system.

```
show ip ospf virtual-link brief
```

Command Mode

- Privileged EXEC
- User EXEC

Display Parameters

Term	Definition
Area ID	The area id of the requested OSPF area.
Neighbor	The neighbor interface of the OSPF virtual interface.
Hello Interval	The configured hello interval for the OSPF virtual interface.
Dead Interval	The configured dead interval for the OSPF virtual interface.
Retransmit Interval	The configured retransmit interval for the OSPF virtual interface.
Transmit Delay	The configured transmit delay for the OSPF virtual interface.

5.3 **Routing Information Protocol Commands**

This section describes the commands you use to view and configure Routing Information Protocol (RIP), which is a distance-vector routing protocol that you use to route traffic within a small network.

5.3.1 **Routing Information Protocol Commands**

5.3.1.1 **router rip**

Use this command to enter Router RIP mode.

```
router rip
```

Command Mode

Global Config

5.3.1.2 **enable (RIP)**

This command resets the default administrative mode of RIP in the router (active). Use the no command to set the administrative mode of RIP in the router to inactive.

```
enable
```

```
no enable
```

Default

Enabled.

Command Mode

Router RIP Config

5.3.1.3 ip rip

This command enables RIP on a router interface or range of interfaces.
Use the no command to disable RIP on a router interface.

```
ip rip
no ip rip
```

Default

Disabled.

Command Mode

Interface Config

5.3.1.4 auto-summary

This command enables the RIP auto-summarization mode.
Use the no command to disable the RIP auto-summarization mode.

```
auto-summary
no auto-summary
```

Default

Disabled.

Command Mode

Router RIP Config

5.3.1.5 default-information originate (RIP)

This command is used to control the advertisement of default routes.
Use the no command to control the advertisement of default routes.

```
default-information originate
no default-information originate
```

Command Mode

Router RIP Config

5.3.1.6 default-metric (RIP)

This command is used to set a default for the metric of distributed routes.
Use the no command to reset the default metric of distributed routes to its default value.

```
default-metric 0-15
no default-metric
```

Command Mode

Router RIP Config

5.3.1.7 distance rip

This command sets the route preference value of RIP in the router. Lower route preference values are preferred when determining the best route. A route with a preference of 255 cannot be used to forward traffic.

Use the no command to set the default route preference value of RIP in the router.

```
distance rip 1-255
no distance rip
```

Command Mode

Router RIP Config

5.3.1.8 **distribute-list out (RIP)**

This command is used to specify the access list to filter routes received from the source protocol.

Use the no command to specify the access list to filter routes received from the source protocol.

```
distribute-list 1-199 out {ospf | bgp | static | connected}
no distribute-list 1-199 out {ospf | bgp | static | connected}
```

Default

0

Command Mode

Router RIP Config

5.3.1.9 **ip rip authentication**

This command sets the RIP Version 2 Authentication Type and Key for the specified interface or range of interfaces. The value of type is either none, simple, or encrypt. The value for authentication key [*key*] must be 16 bytes or less. The [*key*] is composed of standard displayable, noncontrol keystrokes from a Standard 101/102-key keyboard. If the value of type is *encrypt*, a keyid in the range of 0 and 255 must be specified. Unauthenticated interfaces do not need an authentication key or authentication key ID.

Use the no command to set the default RIP Version 2 Authentication Type for an interface.

```
ip rip authentication {none | {simple key} | {encrypt key keyid}}
no ip rip authentication
```

Default

None.

Command Mode

Interface Config

5.3.1.10 **ip rip receive version**

This command configures an interface or range of interfaces to allow RIP control packets of the specified version(s) to be received.

The value for mode is one of: *rip1* to receive only RIP version 1 formatted packets, *rip2* for RIP version 2, *both* to receive packets from either format, or *none* to not allow any RIP control packets to be received.

Use the no command to configure the interface to allow RIP control packets of the default version(s) to be received.

```
ip rip receive version {rip1 | rip2 | both | none}
no ip rip receive version
```

Default

Both.

Command Mode

Interface Config

5.3.1.11 ip rip send version

This command configures an interface or range of interfaces to allow RIP control packets of the specified version to be sent. The value for mode is one of: `rip1` to broadcast RIP version 1 formatted packets, `rip1c` (RIP version 1 compatibility mode) which sends RIP version 2 formatted packets via broadcast, `rip2` for sending RIP version 2 using multicast, or `none` to not allow any RIP control packets to be sent.

Use the `no` command to configure the interface to allow RIP control packets of the default version to be sent.

```
ip rip send version {rip1 | rip1c | rip2 | none}
no ip rip send version
```

Default

rip2.

Command Mode

Interface Config

5.3.1.12 hostroutesaccept

This command enables the RIP hostroutesaccept mode.

Use the `no` command to disable the RIP hostroutesaccept mode.

```
hostroutesaccept
no hostroutesaccept
```

Default

Enabled.

Command Mode

Router RIP Config

5.3.1.13 split-horizon

This command sets the RIP split horizon mode. Split horizon is a technique for avoiding problems caused by including routes in updates sent to the router from which the route was originally learned. The options are: `None` - no special processing for this case. `Simple` - a route will not be included in updates sent to the router from which it was learned. `Poisoned reverse` - a route will be included in updates sent to the router from which it was learned, but the metric will be set to infinity.

Use the `no` command to set the default RIP split horizon mode.

```
split-horizon {none | simple | poison}
no split-horizon
```

Default

simple.

Command Mode

Router RIP Config

5.3.1.14 redistribute (RIP)

This command configures RIP protocol to redistribute routes from the specified source protocol/routers. There are five possible match options. When you submit the command `redistribute ospf match match-type` the matchtype or types specified are added to any match types presently being redistributed. Internal routes are redistributed by default.

Use the `no` command to de-configure RIP protocol to redistribute routes from the specified source protocol/routers.

For OSPF as source protocol

```
redistribute ospf [metric 0-15] [match [internal] [external 1] [external 2] [nssaexternal 1] [nssa-external-2]]
```

For other source protocol

```
redistribute {bgp | static | connected} [metric 0-15]
no redistribute {ospf | bgp | static | connected} [metric]
[match [internal] [external 1] [external 2] [nssa-external 1] [nssa-external-2]]
```

Default

- metric - not-configured
- match - internal

Command Mode

Router RIP Config

5.3.1.15 show ip rip

This command displays information relevant to the RIP router.

```
show ip rip
```

Command Mode

- Privileged EXEC
- User EXEC

Display Parameters

Term	Definition
RIP Admin Mode	Enable or disable.
Split Horizon Mode	None, simple or poison reverse.
Auto Summary Mode	Enable or disable. If enabled, groups of adjacent routes are summarized into single entries, in order to reduce the total number of entries. The default is enable.
Host Routes Accept Mode	Enable or disable. If enabled the router accepts host routes. The default is enable.
Global Route Changes	The number of route changes made to the IP Route Database by RIP. This does not include the refresh of a route's age.
Global queries	The number of responses sent to RIP queries from other systems.
Default Metric	The default metric of redistributed routes if one has already been set, or blank if not configured earlier. The valid values are 1 to 15.
Default Route Advertise	The default route.

5.3.1.16 show ip rip interface brief

This command displays general information for each RIP interface. For this command to display successful results routing must be enabled per interface (i.e., ip rip).

```
show ip rip interface brief
```

Command Mode

- Privileged EXEC
- User EXEC

Display Parameters

Term	Definition
Interface	unit/slot/port
IP Address	The IP source address used by the specified RIP interface.
Send Version	The RIP version(s) used when sending updates on the specified interface. The types are none, RIP-1, RIP-1c, RIP-2.
Receive Version	The RIP version(s) allowed when receiving updates from the specified interface. The types are none, RIP-1, RIP-2, Both.
RIP Mode	The administrative mode of router RIP operation (enabled or disabled).
Link State	The mode of the interface (up or down).

5.3.1.17 show ip rip interface

This command displays information related to a particular RIP interface. The argument *unit/slot/port* corresponds to a physical routing interface or VLAN routing interface. The keyword *vlan* is used to specify the VLAN ID of the routing VLAN directly instead of a unit/slot/port format.

```
show ip rip interface {unit/slot/port | vlan 1-4093}
```

Command Mode

- Privileged EXEC
- User EXEC

Display Parameters

Term	Definition
Interface	<i>unit/slot/port</i> . This is a configured value.
IP Address	The IP source address used by the specified RIP interface. This is a configured value.
Send Version	The RIP version(s) used when sending updates on the specified interface. The types are none, RIP-1, RIP-1c, RIP-2. This is a configured value.
Receive Version	The RIP version(s) allowed when receiving updates from the specified interface. The types are none, RIP-1, RIP-2, Both. This is a configured value.
RIP Admin Mode	RIP administrative mode of router RIP operation; enable activates, disable de-activates it. This is a configured value.
Link State	Indicates whether the RIP interface is up or down. This is a configured value.
Authentication Type	The RIP Authentication Type for the specified interface. The types are none, simple, and encrypt. This is a configured value.

The following information will be invalid if the link state is down.

Term	Definition
Bad Packets Received	The number of RIP response packets received by the RIP process which were subsequently discarded for any reason.
Bad Routes Received	The number of routes contained in valid RIP packets that were ignored for any reason.
Updates Sent	The number of triggered RIP updates actually sent on this interface.

Chapter 6

Troubleshooting

6.1 Troubleshooting

- Verify that the device is using the right power cord/adaptor (DC 24-110V); please do not use a power adaptor with DC output higher than 110V, or the device may be damaged.
- Select the proper UTP/STP cable to construct the user network. Use unshielded twisted-pair (UTP) or shield twisted-pair (STP) cable for M12 connections that depend on the connector type the switch equipped: 100R Category 3, 4 or 5 cable for 10Mbps connections, 100R Category 5 cable for 100Mbps connections, or 100R Category 5e/above cable for 1000Mbps connections. Also ensure that the length of any twisted-pair connection does not exceed 100 meters (328 feet).
R = replacement letter for Ohm symbol.
- **Diagnosing LED Indicators:** To assist in identifying problems, the switch can be easily monitored through panel indicators, which describe common problems the user may encounter, so the user can be guided towards possible solutions.
- If the power indicator does not light on when the power cord is plugged in, you may have a problem with power cord. Check for loose power connections, power losses, or surges, at the power outlet. If you still cannot resolve the problem, contact a local dealer for assistance.
- If the LED indicators are normal and the connected cables are correct but packets still cannot be transmitted, please check the user system's Ethernet device configuration or status.

ADVANTECH

Enabling an Intelligent Planet

www.advantech.com

Please verify specifications before quoting. This guide is intended for reference purposes only.

All product specifications are subject to change without notice.

No part of this publication may be reproduced in any form or by any means, electronic, photocopying, recording or otherwise, without prior written permission of the publisher.

All brand and product names are trademarks or registered trademarks of their respective companies.

© Advantech Co., Ltd. 2016