VOLKTEK

L3 Fully Managed Ethernet Switch

User Manual

Version 1.2
16 June 2026

VOLKTEK

Revision History

Version	Date	Firmware	Author	Description
0.1	12 Nov 2025	—	Volktek	Initial draft
1.0	4 Mar 2026	V2.0.0.B5	Volktek	First release (firmware V2.0.0.B5)
1.1	9 Mar 2026	V2.0.0.S0	Volktek	Add Dashboard section (3.3); update web timeout range to 0-3600 (4.12.1.8); add DI1/DI2/DO LED indicators to Chapter 1
1.2	16 Jun 2026	V2.1.0.S0	Volktek	Update Web GUI Dashboard and Front Panel; add first-login password change, enforced HTTPS behavior, alarm escalation and alarm monitoring, PoE power clarification; remove obsolete LED Status page.

FCC WARNING

This equipment has been tested and found to comply with the limits for a class A device, pursuant to part 15 of FCC rules. These limits are designed to provide reasonable protection against harmful interference in a commercial installation. This equipment generates uses and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communication. 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 the user's own expense.

CE

This is a Class A product. In a domestic environment, this product may cause radio interference in which case the user may be required to take adequate measures.

HOT& COLD WARNING

The switch surface will get very hot or cold depending on the operating environment. Please take special care when touching the working switch.

GENERAL WARNINGS – READ BEFORE USE

Take special care to read and understand all the content in the warning boxes.

Warning

Do not work on the system or connect or disconnect cables during periods of lightning activity.

Warning

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

Warning

Do not stack the chassis on any other equipment. If the chassis falls, it can cause severe bodily injury and equipment damage.

Warning

An exposed wire lead from a DC-input power source can conduct harmful levels of electricity. Be sure that no exposed portion of the DC-input power source wire extends from the terminal block plug.

Warning

Ethernet cables must be shielded when used in a central office environment.

Warning

If a redundant power system (RPS) is not connected to the switch, install an RPS connector cover on the back of the switch.

Warning

Read the wall-mounting instructions carefully before beginning installation. Failure to use the correct hardware or to follow the correct procedures could result in a hazardous situation for people and damage to the system.

Warning

Before performing any of the following procedures, ensure that power is removed from the DC circuit.

Warning

Read the installation instructions before connecting the system to the power source.

Warning

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

Warning

To prevent bodily injury when mounting or servicing this unit in a rack, you must take special precautions to ensure that the system remains stable. The following guidelines are provided to ensure your safety:

- ✓ **This unit should be mounted at the bottom of the rack if it is the only unit in the rack.**
- ✓ **When mounting this unit in a partially filled rack, load the rack from the bottom to the top with the heaviest component at the bottom of the rack.**
- ✓ **If the rack is provided with stabilizing devices, install the stabilizers before mounting or servicing the unit in the rack.**

Warning

Only trained and qualified personnel should be allowed to install, replace, or service this equipment.

Warning

When installing or replacing the unit, the ground connection must always be made first and disconnected last.

Warning

Voltages that present a shock hazard may exist on Power over Ethernet (PoE) circuits if interconnections are made using uninsulated exposed metal contacts, conductors, or terminals. Avoid using such interconnection methods, unless the exposed metal parts are located within a restricted access location, and users and service people who are authorized within the restricted access location are made aware of the hazard. A restricted access area can be accessed only through the use of a special tool, lock and key or other means of security.

Warning

No user-serviceable parts inside. Do not open.

Warning

This equipment must be grounded. Never disconnect the ground conductor or operate the equipment in the absence of a suitable installed ground conductor. Contact the appropriate electrical inspection authority or an electrician if you are uncertain that suitable grounding is available.

1 Product Overview

This User Manual applies to multiple fully managed Ethernet switch models with advanced layer 3 functionalities. The following sections list the individual models as well as their capabilities and specifications.

1.1 9561-8GP4XS-TSN

The 9561-8GP4XS-TSN Industrial Ethernet Switch is an 8-port 10/100/1000 BASE-T RJ45 and 4-slot FX/GbE/10G SFP+ TSN-capable PoE++ switch with one RJ45 console port and one USB port for firmware and configuration maintenance. This switch is engineered and designed especially for harsh & extreme industrial applications, featuring high tolerance to vibration and shock, and a wide operating temperature range from -40°C to 70°C.

It supports advanced Power over Ethernet capabilities including BT PoE (IEEE 802.3bt) delivering up to 90W, enabling power-hungry devices such as high-end IP cameras and wireless access points to be operated without additional power sources. The PoE management features allow per-port power monitoring, prioritization, and scheduling to optimize power distribution and efficiency.

Additionally, the switch supports precision timing protocols such as IEEE 1588 and 802.1ASRev, making it ideal for applications requiring synchronized operations. Advanced features like G.8032 and full TSN support, along with TSN Domain Protection, enhance network reliability and security.

Interface

Ports	
10/100/1000Base-T (RJ45)	8
10G SFP+ slots	4
Console port (RJ45 to RS232)	1
USB port	1
Digital Inputs	2 x Pairs Dry Contact
Digital Output	1 x Alarm Relay with current carrying capacity of 1 A @ 24 VDC

Power

Primary input voltage	48~57 VDC
Redundant input voltage	48~57 VDC
System Power Consumption	19 W

PoE

Max Power delivered by PSE	90 W
Voltage Range (at PSE)	48-57 V
Output capacity for PoE	360 W

Operating Requirement

Operating temperature	-40□ to 70□ (-40°F~158°F)
Storage temperature	-40□ to 85□ (-40°F~185°F)
Operating humidity	5% to 95% RH (Non-Condensing)
Storage humidity	5% to 95% RH (Non-Condensing)

LED Indicators

This switch is equipped with Unit LEDs to enable you to determine the status of the switch, as well as Port LEDs to display what is happening in all your connections. They are as follows:

Unit LEDs		
LED	**Condition**	**Status**
PWR (Green)	Illuminated	Primary power on
	Off	Primary power failure or not available
RPS (Green)	Illuminated	Redundant power is supplied to switch
	Off	Redundant power off or failure
ALM (Red)	Illuminated	Alarm triggered for abnormal power status and Anomalous features
	Off	Normal operation or DIP switch off
POST (Green)	Illuminated	Switch is ready or running
	Blinking	Self-testing the device when power on
	Off	Switch is not ready
SFP (9~12) (SFP+)	Green	Link speed at 10G
	Amber	Link speed at 1G
	Off	Port disconnected or link failed
PoE (1~8) (Green)	Solid	Supplying power to a powered device (PD)
	Off	PoE is disabled or no PD connected
1000 (1~8) (Green)	Illuminated	Copper port link speed at 1000Mbps
	Off	Copper Port link speed at 10/100Mbps
LNK/ACT (Green)	Illuminated	port link up
	Blinking	Activity (receiving or transmitting data)
	Off	Port disconnected or link failed

DIP Switches

Power: DIP 1 and DIP 2 is for primary power and redundant power supply.

No	Name	Description
1	PWR	ON : Primary power alarm reporting is enabled OFF: Primary power alarm reporting is disabled
2	RPS	ON : Redundant power alarm reporting is enabled OFF: Redundant power alarm reporting is disabled

Console Port

- Connect your computer to the console port on the switch using the appropriate cable.
- Use terminal emulation software with the following settings:

Default Settings for the Console Port

Setting	Default Value

Terminal Emulation	VT100
Baud Rate	115200
Parity	None
Number of Data Bits	8
Number of Stop Bits	1
Flow Control	None

- Press [ENTER] to open the login screen.

Setting	Default Value
Default Username	admin
Default Password	admin

1.2 9561-8GT4XS-TSN

The 9561-8GT4XS-TSN industrial network switch is designed to deliver high performance and reliability in demanding industrial environments. It features a versatile interface with 8x 10/100/1000Mbps RJ45 ports and 4x FX/GbE/10G SFP+ ports, ensuring robust and flexible connectivity options.

The switch supports precision timing protocols such as IEEE 1588 and 802.1AS-Rev, making it ideal for applications requiring synchronized operations. Advanced features like G.8032 and full TSN support, along with TSN Domain Protection (802.1Qci, 802.1Qav, 802.1Qbv, 802.1Qbu/802.1br), enhance network reliability and security. High traffic availability is maintained through standards like 802.1CB, IEC- 62439-2 2016 MRP, and HSR/PRP, while time distribution capabilities include 802.1AS-2020 (gPTP), IEEE1588:2019, and 1PPS In/Out. Efficient power management is ensured with 802.3az (EEE) and Wake-on-LAN (WoL) capabilities.

Interface

Ports	
10/100/1000Base-T (RJ45)	8
10G SFP+ slots	4
Console port (RJ45 to RS232)	1
USB port	1
Digital Inputs	2 x Pairs Dry Contact
Digital Output	1 x Alarm Relay with current carrying capacity of 1 A @ 24 VDC

Power

Primary input voltage	12~48 VDC
Redundant input voltage	12~48 VDC
System Power Consumption	19 W

Operating Requirement

Operating temperature	-40□ to 70□ (-40°F~158°F)
Storage temperature	-40□ to 85□ (-40°F~185°F)
Operating humidity	5% to 95% RH (Non-Condensing)
Storage humidity	5% to 95% RH (Non-Condensing)

LED Indicators

This switch is equipped with Unit LEDs to enable you to determine the status of the switch, as well as Port LEDs to display what is happening in all your connections. They are as follows:

Unit LEDs		
LED	**Condition**	**Status**

PWR (Green)	Illuminated	Primary power on
	Off	Primary power failure or not available
RPS (Green)	Illuminated	Redundant power is supplied to switch
	Off	Redundant power off or failure
ALM (Red)	Illuminated	Alarm triggered for abnormal power status and Anomalous features
	Off	Normal operation or DIP switch off
POST (Green)	Illuminated	Switch is ready or running
	Blinking	Self-testing the device when power on
	Off	Switch is not ready
SFP (9~12) (SFP+)	Green	Link speed at 10G
	Amber	Link speed at 1G
	Off	Port disconnected or link failed
1000 (1~8) (Green)	Illuminated	Copper port link speed at 1000Mbps
	Off	Copper Port link speed at 10/100Mbps
LNK/ACT (Green)	Illuminated	port link up
	Blinking	Activity (receiving or transmitting data)
	Off	Port disconnected or link failed

DIP Switches

Power: DIP 1 and DIP 2 is for primary power and redundant power supply.

No	Name	Description
1	PWR	ON : Primary power alarm reporting is enabled OFF: Primary power alarm reporting is disabled
2	RPS	ON : Redundant power alarm reporting is enabled OFF: Redundant power alarm reporting is disabled

Console Port

- Connect your computer to the console port on the switch using the appropriate cable.
- Use terminal emulation software with the following settings:

Default Settings for the Console Port

Setting	Default Value
Terminal Emulation	VT100
Baud Rate	115200
Parity	None
Number of Data Bits	8
Number of Stop Bits	1
Flow Control	None

- Press [ENTER] to open the login screen.

Setting	Default Value
Default Username	admin

Default Password	admin

2 About the Switch

This section contains general information applying to all switch models as well as a software feature table, indicating which features are supported by each switch model.

2.1 Software Feature Overview

	9561-8GP4XS-TSN	9561-8GT4XS-TSN				
Management						
CLI	Y	Y				
Telnet/SSHv2	Y	Y				
HTTP/HTTPs	Y	Y				
Syslog	Y	Y				
NTPv4 Client	Y	Y				
DHCP Server	Y	Y				
DHCP Client	Y	Y				
Server Control	Y	Y				
SFP Info	Y	Y				
SFP DDMI	Y	Y				
Cable Test	Y	Y				
Hardware Monitor	-	-				
Port Utilization	Y	Y				
Alarm Info	Y	Y				
Port Management	Y	Y				
Remote Reboot	Y	Y				
Account authority	Y	Y				
MAC Aging Time	Y	Y				
User Self-defined Configuration	Y	Y				
Alarm LED Event Trigger	-	-				
EEE	Y	Y				
USB Management	-	-				
DNS Client/Relay	Y	Y				
IPv4 Traceroute	Y	Y				
FTP Client	Y	Y				
SNMP v1/v2C/v3	Y	Y				
SNMP Trap	Y	Y				
SNMP Infom	Y	Y				
Management VLAN	Y	Y				
LLDP	Y	Y				
LLDP-MED	Y	Y				
UDLD	Y	Y				
DHCP Relay	Y	Y				
DHCP Option82	Y	Y				
SPAN	Y	Y				
RSPAN	Y	Y				
VSPAN	Y	Y				
Rule-based Flow Mirroring	Y	Y				
RMON (1,2,3,9)	Y	Y				
Standard MIBs	Y	Y				
Private MIBs	Y	Y				
sFlow	Y	Y				

802.3ah EFM	Y	Y				
802.3ag CFM	Y	Y				
Y.1731	Y	Y				
Synchronization						
PTP	Y	Y				
IEEE 802.1AS	Y	Y				
IEEE 802.1CB FRER	Y	Y				
Reliability						
STP/RSTP/MSTP	Y	Y				
ERPs v1/v2	Y	Y				
Link Aggregation	Y	Y				
Dual Image	Y	Y				
G.8031 APS	Y	Y				
IEC 62439-2 MRP	Y	Y				
IEC 62439-3 PRP	Y	Y				
IEC 62439-3 HSR	Y	Y				
VLAN						
802.1Q VLAN	Y	Y				
GARP/GVRP	Y	Y				
MVRP	Y	Y				
Port Isolation	Y	Y				
Private VLAN	Y	Y				
MAC-Based VLAN	Y	Y				
IP Subnet-Based VLAN	Y	Y				
Protocol-Based VLAN	Y	Y				
QinQ	Y	Y				
VLAN Translation	Y	Y				
Voice VLAN	Y	Y				
VLAN Trunking	Y	Y				
iPVLAN	Y	Y				
Traffic Control						
IGMP Snooping v1/v2/v3	Y	Y				
MVR	Y	Y				
802.1ak MRP	Y	Y				
802.1p QoS	Y	Y				
802.1Qbv	Y	Y				
802.1Qbu & 802.3br	Y	Y				
802.1Qci	Y	Y				
Policy	Y	Y				
Shaper	Y	Y				
Storm Control	Y	Y				
Flow Control	Y	Y				
Loop Detection	Y	Y				
L3 Routing						
Inter-VLAN Routing	Y	Y				
Static Route	Y	Y				
RIPv2	Y	Y				
OSPFv2	Y	Y				
Security						
DHCP Snooping	Y	Y				
IP Source Guard	Y	Y				
ARP Inspection	Y	Y				
ACL	Y	Y				
TACACS+	Y	Y				

RADIUS	Y	Y				
802.1X Authentication (Port-Based / MAC)	Y	Y				
Port Security	Y	Y				
Sticky MAC	Y	Y				
BPDU Guard/Filter	Y	Y				
Root Guard	Y	Y				
Managed Host	Y	Y				
MAC Anti-Spoofing	Y	Y				
DoS/DDoS	Y	Y				
SCP/SFTP	Y	Y				
Secure Boot	Y	Y				
IPv6						
IPv6 Telnet/SSH	Y	Y				
IPv6 HTTP/HTTPs	Y	Y				
NTPv6 Client	Y	Y				
DHCPv6 Client	Y	Y				
DHCPv6 Shield	Y	Y				
IPv6 Default Gateway	Y	Y				
IPv6 ARP (Neighbor Discovery)	Y	Y				
IPv6 SNMPv1/v2c/v3	Y	Y				
IPV6 SNMP Trap	Y	Y				
DHCPv6 Relay	Y	Y				
DHCPv6 Option 18	Y	Y				
IPv6 MLDv1/v2	Y	Y				
IPv6 Static Route	Y	Y				
IPv6 OSPFv3	Y	Y				
IPv6 Inter-VLAN Routing	Y	Y				
IPv6 ACL	Y	Y				
IPv6 TACACS+	Y	Y				
IPv6 RADIUS	Y	Y				
IPv6 Traceroute	Y	Y				
IPv6 Source Guard	Y	Y				
PoE						
PoE Power On/Off	Y	-				
PoE Priority	Y	-				
Power Budget Control Per (System/Port)	Y	-				
Power Class Detection	Y	-				

2.2 Hardware Installation

The location chosen for installing the switch may greatly affect its performance. When selecting, we recommend considering the following rules:

- ✓ Install the switch in an appropriate place. See Product Overview for the acceptable temperature and humidity ranges.
- ✓ Install the switch in a location that is not affected by strong electromagnetic field generators (such as motors), vibration, dust, and direct sunlight.
- ✓ Leave at least 10cm of space at the front and rear of the unit for ventilation.

ATTENTION

The switch is an open-type device and shall be mounted on a DIN rail or wall-mounted inside a cabinet or enclosure

Hardware Installation

- ✓ **Step 1**: Unpack the device and other contents of the package.
- ✓ **Step 2**: Fasten DIN-Rail or Wall-mount kit on the rear of the switch
- ✓ **Step 3:** Connect the power to the power terminal block.
- ✓ **Step 4**: Connect the Ethernet (RJ45) port to the networking device and check the LED status to confirm the connection is established.

DIN Rail Installation

The switch has a DIN rail bracket on the back of the switch and can be DIN-Rail-mounted in cabinet or enclosure.

Mounting the Switch

Place the switch on the DIN rail from above using the slot. Push the front of the switch toward the mounting surface until it snaps into place with a click sound

Dismounting the Switch

Pull out the lower edge of the switch and then remove the switch from the DIN rail.

Grounding the Switch

Before powering the switch, ground the switch to earth.
Ensure the rack on which the switch is to be mounted is properly grounded and incompliance with ETSI ETS 300 253. Verify that there is a good electrical connection to the grounding point on the rack (no paint or isolating surface treatment).

ATTENTION

This product is intended to be mounted to a well-grounded mounting surface such as a metal panel.

CAUTION

The earth connection must not be removed unless all power supply connections have been disconnected.
The device is installed in a restricted-access location; it has a separate protective earthing terminal on the chassis that must be permanently connected to earth ground to adequately ground the chassis and protect the operator from electrical hazards.

ATTENTION

The product should be mounted in an Industrial Control Panel, and the ambient temperature should not exceed 75°C.

ATTENTION

A corrosion-free mounting rail is advisable.
When installing, make sure to allow for enough space to properly install the cabling.

WARNING

- Wiring cable temperature should support at least 105°C
- Tighten the wire to a torque value 5lb

- Use the following wire gauge range for the terminal block:

Non-PoE	PoE (240W)	PoE (360 W delivery budget; 400 W PSU peak)
12~24 AWG	12~18 AWG	12~16 AWG

WARNING

Safety measures should be taken before connecting the power cable**.** Turn off the power before connecting modules or wires. The correct power supply voltage is listed on the product label. Check the voltage of your power source to make sure that you are using the correct voltage. **DO NOT** use a voltage greater than what is specified on the product label. Calculate the maximum possible current in each power wire and common wire. Observe all electrical codes dictating the maximum current allowable for each wire size. If current exceeds the maximum rating, the wiring can overheat causing serious damage to your equipment.

Please read and follow these guidelines:

- Use separate paths to route wiring for power and devices. If power wiring and device wiring paths must cross, make sure the wires are perpendicular to the intersection point.
  Note: Do not run signal or communications wiring and power wiring through the same wire conduit. To avoid interference, wires with different signal characteristics should be routed separately.
- You can use the type of signal transmitted through a wire to determine which wires should be kept separate. The rule of thumb is that wiring that shares similar electrical characteristics can be bundled together
- You should separate input wiring from output wiring
- We advise that you label the wiring to all devices in the system.

Wiring the Alarm Contact

The Alarm Contact consists of the two last contacts of the terminal block on the switch's top panel labeled **ALM**. They are used to detect both power faults and port faults. The two wires attached to the ALM contacts form an open circuit when the switch has lost power from one of the DC power inputs. Otherwise, the alarm circuit will be closed.

Reset Button

There is a *Reset* button in front of switch which can help to manually hardware reboot.

2.3 Interface Connectors

The switch utilizes ports with copper and SFP/SFP+ fiber port connectors functioning under Ethernet/Fast Ethernet/Gigabit/10 Gigabit Ethernet standards.

10/100/1000Base-T Ports

The 10/100/1000Base-T ports support network speeds of 10Mbps, 100Mbps or 1000Mbps, and can operate in half- and full-duplex transfer modes. These ports also offer automatic MDI/MDI-X crossover detection that gives true "plug-n-play" capability – just plug the network cables into the ports and the ports will adjust according to the end-node devices. The following are recommended cabling for the RJ45 connectors: (1) 10Mbps – Cat 3 or better; (2) 100/1000Mbps – Cat 5e or better.

SFP Slots for SFP/SFP+ Modules

The four SFP slots are designed to Gigabit SFP/10 Gigabit SFP+ modules that support network speed of 1000Mbps/10Gbps.

Installing the SFP/SFP+ Modules and Fiber Cable

1. Slide the selected SFP/SFP+ module into the selected SFP slot (Make sure the SFP/SFP+ module is aligned correctly with the inside of the slot)
2. Insert and slide the module into the SFP slot until it clicks into place
3. Remove any rubber plugs that may be present in the SFP/SFP+ module's mouth
4. Align the fiber cable's connector with the SFP/SFP+ module's mouth and insert the connector
5. Slide the connector in until a click is heard
6. If you want to pull the connector out, first push down the release clip on top of the connector to release the connector from the SFP/SFP+ module.

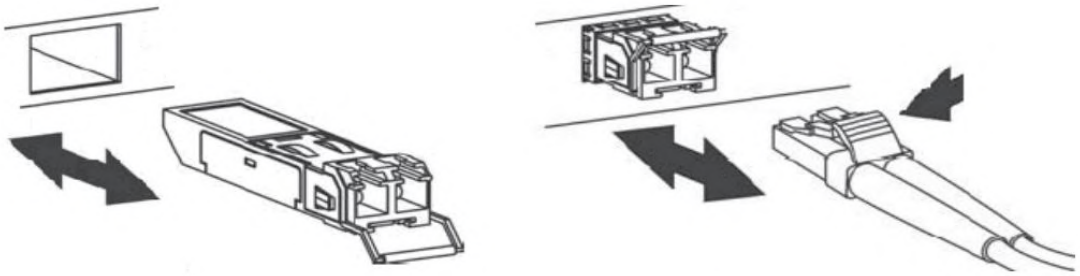

To properly connect fiber cabling: Check that the fiber terminators are clean. You can clean the cable plugs by wiping them gently with a clean tissue or cotton ball moistened with a little ethanol. Dirty fiber terminators on fiber optic cables will impair the quality of the light transmitted through the cable and lead to degraded performance on the port.

Note: *When inserting the cable, be sure the tab on the plug clicks into position to ensure that it is properly seated.*
Check the corresponding port of LED on the switch to ensure that the connection is valid. (Refer to the LED chart).

2.4 Entering the CLI

Press the [Enter] key to enter the login command prompt when the message below is displayed on the screen.
`Press Enter to get started`

Input "**`admin`**" to enter the CLI mode when the message below is displayed on the screen.
`Username:`

You can execute a few limited commands depending on the current user privilege when the CLI prompt is displayed as below.
`#`

If you want to execute more powerful commands or change the configuration, you must enter the configuration mode.
Input command "**`configure terminal`**"
`(config)#`

3 Web GUI

Volktek managed Ethernet switches offer a comprehensive and intuitive Web-based interface to configure management features as well as monitor switch information and perform maintenance tasks.
The web-based graphical user interface (Web GUI) is accessible with a web browser at the IP address assigned to the switch.

3.1 Login

The switch can be managed through the Web GUI with a standard web browser. Connect the management PC to the switch management network and open the switch IP address. The factory default IP address is 192.168.0.254.
The factory default user name is admin and the factory default password is admin.
For firmware V2.1.0.S0, the first login with the default password opens the Default Password Must Be Changed page before entering the Web GUI. Enter a new password and confirm it, then click Change Password. Click Skip to enter the Web GUI without changing the password; the prompt appears again on the next default-password login.
The new password must use at least eight characters and include uppercase letters, lowercase letters, digits, and special characters. The New Password and Confirm Password fields must match before the password can be changed.
If repeated login attempts fail, the login service may temporarily delay or block additional attempts. Wait for the lockout period to expire, or use a valid administrator account.
Figure: First login default-password change page captured from the Web GUI.

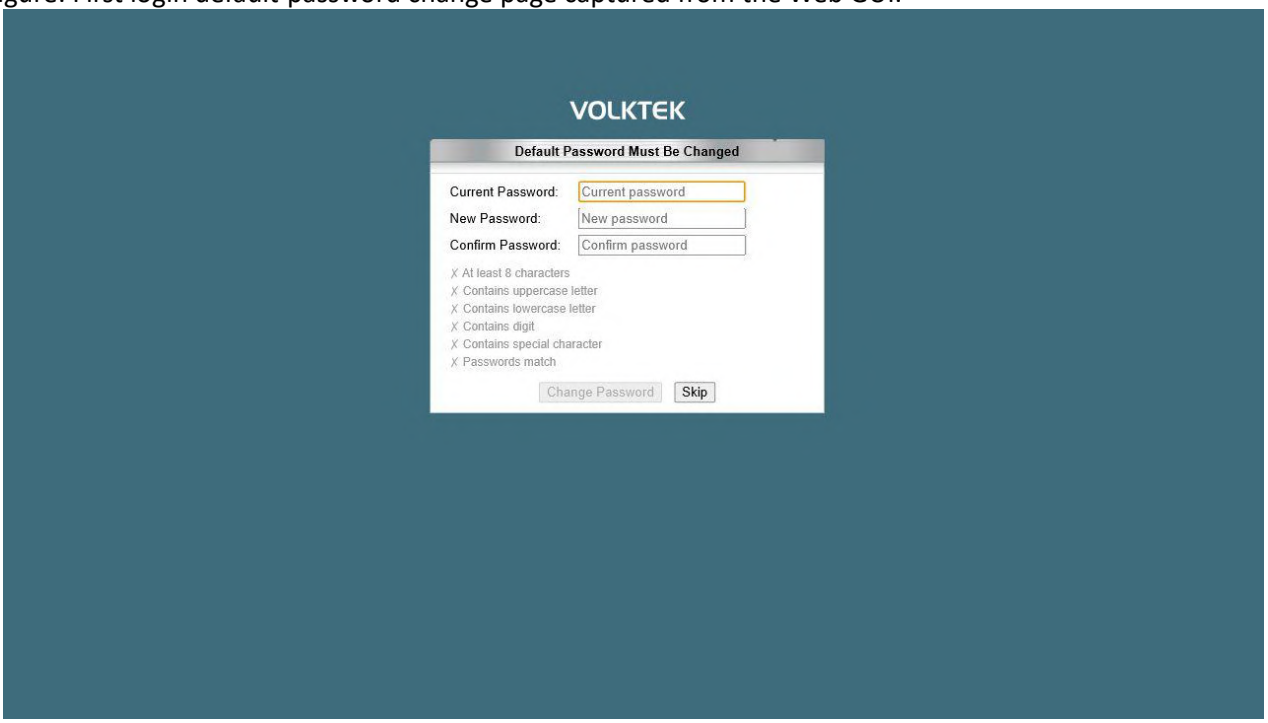

3.2 Navigation

The feature menu can be navigated using the sidebar on the left side of the screen, which is separated into five multi-level drop-down sections.

- The Configuration section contains all settings that can be used to adjust features of the switch.
- The Monitor section shows status information about the switch. Most configuration pages will have a corresponding monitor page to observe the effects of the settings.
- The Diagnostics section contains tools to troubleshoot connectivity issues.
- The Maintenance section is meant for administrators to perform changes on configuration files or firmware.
- The Volktek Support Section contains information about how to get in contact with Volktek technical support.

On the top right of the Web GUI are three buttons. The left button will lead you back to the landing page, the center button can be used to log out of the Web GUI, and the right button offers helpful information about settings and parameters for the current page.

3.3 Dashboard

The Dashboard page is shown in the main frame after login. It provides live monitoring for system resources, port traffic, PTP status, and ping targets.
The Dashboard controls are separated into polling, sampling, display window, and session behavior. The parameter table at the end of this section summarizes these settings.
The left panel lists available observables. System includes CPU (All) and Memory. Port 1 through Port 12 each include Rate, Util, Err, Drop, and LinkDown. PTP includes PTP 0 through PTP 3 and the corresponding Pdelay entries. Ping provides + Add Ping Target.
The right View panel is the dashboard workspace. Drag observables from the left panel into the View panel, or double-click an observable to add it.
Figure: Dashboard page captured from the Web GUI.

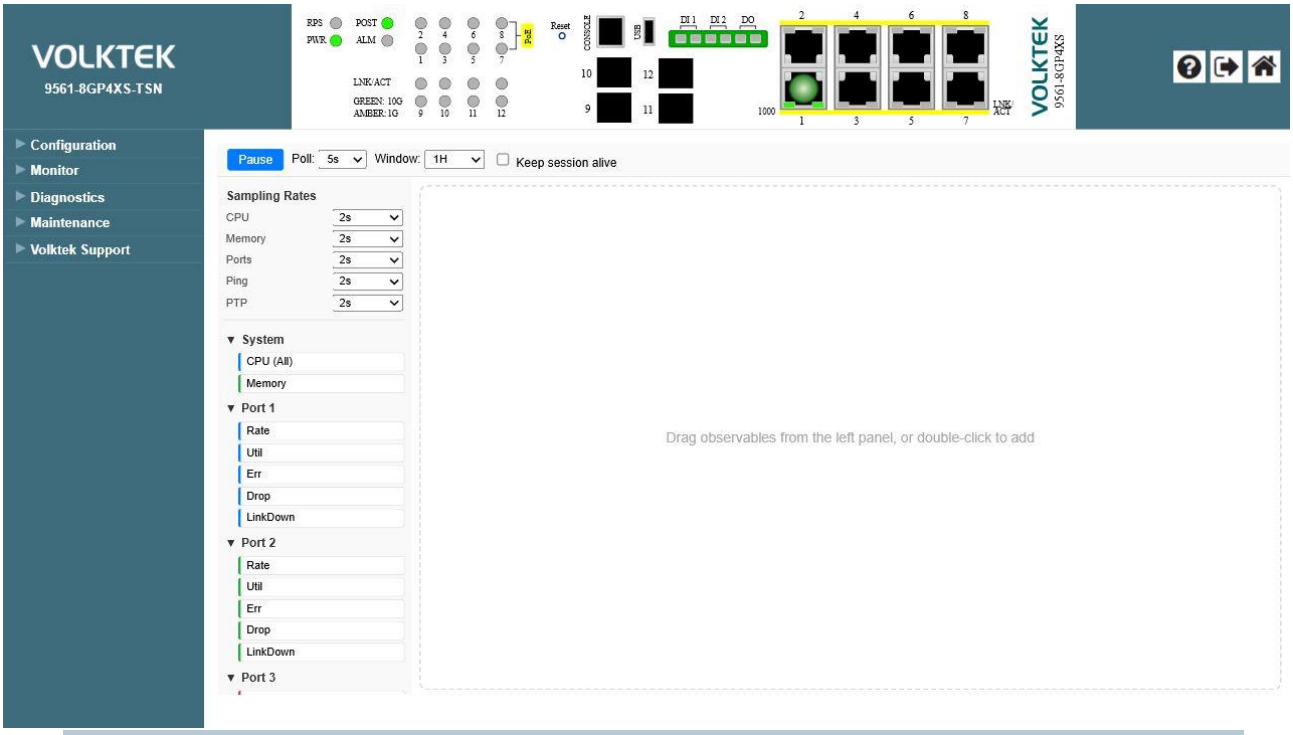

Parameter	Description
Poll interval	Controls how often the Dashboard refreshes data from the switch Web GUI. Available values: 1s, 5s, 10s.
Sampling interval	Controls the per-observable sampling interval for CPU, Memory, Ports, Ping, and PTP data collection. Sampling Rates are independent from the Poll interval. Available values: 100ms, 500ms, 1s, 2s, 5s, and 10s for each observable group.
Window scale	Controls the time range displayed in the Dashboard view. Available values: 30Min, 1H, 3H, 6H, 12H, and 24H.
Keep session alive	Keeps the Web GUI session active while the Dashboard remains open. Unchecked: normal session timeout applies. Checked: Dashboard keeps the session alive.

3.4 Front Panel

The Web GUI Front Panel is a live SVG banner displayed at the top of Web GUI pages. It shows the current hardware state of the switch.

The front panel shows system LEDs for PWR, RPS, ALM, and POST; copper port link/activity/speed state for ports 1 through 8; SFP/SFP+ port state for ports 9 through 12; PoE delivery status for PoE-capable copper ports; and DI1, DI2, and DO status.

The Front Panel is updated by the Web GUI monitoring service and can be used as a quick visual indication of link, power, alarm, PoE, and digital I/O state.

Figure: V2.1.0.S0 live Web GUI Front Panel banner captured from 9561-8GP4XS-TSN.

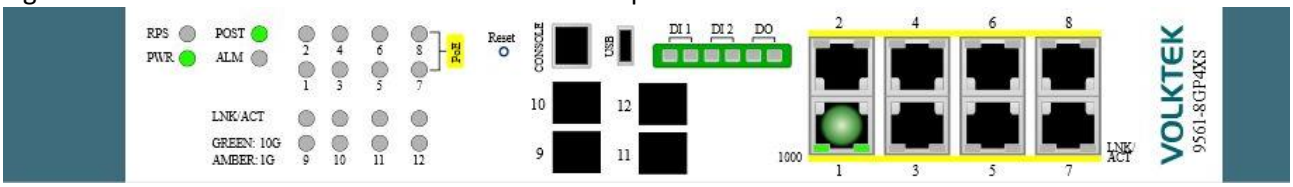

The visual components on the SVG front panel are interactive. Double-click a component, such as a port, to open a detail widget for that component. For a port, the widget can show link state, speed, packet counters, rates, error counters, drop counters, and LLDP peer information when available.

Left-click the popup window to close it.

Figure: Front Panel component detail widget opened by double-clicking a port component.

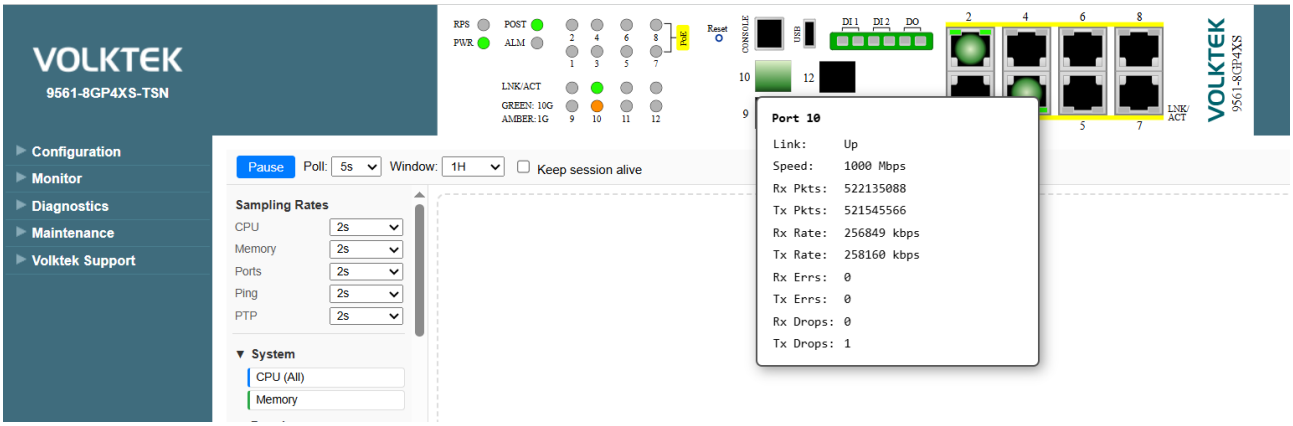

For detailed counter values or event history, use the corresponding Monitor pages such as Port Statistics, PoE, System Log, and Detailed Log.

4 Configuration

The *Configuration* menu contains all the settings that can be made for the management features of the Ethernet switch. Most features have a corresponding tab in the Monitor section located further down in the sidebar, where the configuration changes, and the status of the switch can be observed.
Whenever changes are made, they need to be applied using the *Save* button. The *Reset* button undoes any changes made locally and reverts to previously saved values.

When navigating tables, the >> button will use the last entry of the currently displayed table as a basis for the next lookup. When the end is reached the text *No more entries* are shown in the displayed table. Use the |<< button to start over.

Some configuration tables include clickable icons to modify the table rows. You can modify entry in the table using the following buttons:
: Edits the entry.
: Deletes the entry.
: Adds a new entry.

4.1 System

The *System* settings contain general settings like system information, IP address assignment, time settings for logging and monitoring settings with SNMP.

4.1.1 Information

There are three fields in this System Information Configuration webpage: System Contact, System Name, and System Location, which can also be viewed at Monitor→System→Information.

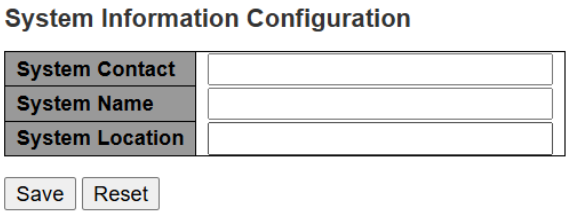

Parameter	Description
System Contact	The textual identification of the contact person for this managed node, together with information on how to contact this person. The allowed string length is 0 to 255, and the allowed content is the ASCII characters from 32 to 126.
System Name	An administratively assigned name for this managed node. By convention, this is the node's fully-qualified domain name. A domain name is a text string drawn from the alphabet (A-Za-z), digits (0-9), minus sign (-). No space characters are permitted as part of a name. The first character must be an alpha character. And the first or last character must not be a minus sign. The allowed string length is 0 to 255.
System Location	The physical location of this node (e.g., telephone closet, 3rd floor). The allowed string length is 0 to 255, and the allowed content is the ASCII characters from 32 to 126.

In this subsection, you can configure the IP settings of the managed switch. It is divided into three parts: IP Configuration, IP Interfaces, and IP Routes. The IP Configuration section defines how the switch operates as an IP host. The IP Status including interfaces and routes can be viewed under Monitor→System→IP Status.

IP Configuration

Domain Name	No Domain Name	
Mode	Host	
DNS Server 0	Configured IPv4 or IPv6	8.8.8.8
DNS Server 1	No DNS server	
DNS Server 2	No DNS server	
DNS Proxy	☐	

Parameter	Description
Domain Name	The name string of local domain where the device belongs. Most queries for names within this domain can use short names relative to the local domain. The system then appends the domain name as a suffix to unqualified names. For example, if domain name is set as 'example.com' and you specify the PING destination by the unqualified name as 'test', then the system will qualify the name to be 'test.example.com'. The following modes are supported: **`No Domain Name`**: No domain name will be used. **`Configured Domain Name`**: Explicitly specify the name of local domain. Make sure the configured domain name meets your organization's given domain. **`From any DHCPv6 interfaces`**: The first domain name offered from a DHCPv6 lease to a DHCPv6-enabled interface will be used. **`From this DHCPv6 interface`**: Specify from which DHCPv6-enabled interface a provided domain name should be preferred.
Mode	Configure whether the IP stack should act as a **`Host`** or a **`Router`**. In **`Host`** mode, IP traffic between interfaces will not be routed. In **`Router`** mode traffic is routed between all interfaces.
DNS Server	This setting controls the DNS name resolution done by the switch. There are four servers available for configuration, and the index of the server presents the preference (less index has higher priority) in doing DNS name resolution. The following modes are supported: **`No DNS server`**: No DNS server will be used. **`Configured IPv4`**: Explicitly provide the valid IPv4 unicast address of the DNS Server in dotted decimal notation. Make sure the configured DNS server could be reachable (e.g. via PING) for activating DNS service. **`Configured IPv6`**: Explicitly provide the valid IPv6 unicast (except linklocal) address of the DNS Server. Make sure the configured DNS server could be reachable (e.g. via PING6) for activating DNS service. **`From any DHCPv4 interfaces:`** The first DNS server offered from a DHCPv4 lease to a DHCPv4-enabled interface will be used. **`From this DHCPv4 interface`**: Specify from which DHCPv4-enabled interface a provided DNS server should be preferred. **`From any DHCPv6 interfaces`**: The first DNS server offered from a DHCPv6 lease to a DHCPv6-enabled interface will be used. **`From this DHCPv6 interface`**: Specify from which DHCPv6-enabled interface a provided DNS server should be preferred.
DNS Proxy	When DNS proxy is enabled, system will relay DNS requests to the currently configured DNS server, and reply as a DNS resolver to the client devices on the network. Only IPv4 DNS proxy is now supported.

The *IP Interfaces* section allows you to configure IPv4 and IPv6 addresses and DHCP settings. A maximum of 126 interfaces is supported.

IP Interfaces

Delete	IF	DHCPv4							
		Enable	Client ID				Hostname	Fallback	Current Lease
			Type	IfMac	ASCII	HEX			
☐	VLAN 1	☐	Auto	Port 1				0	

IPv4		DHCPv6			IPv6	
Address	Mask Length	Enable	Rapid Commit	Current Lease	Address	Mask Length
192.168.203.160	24	☐	☐			

Parameter	Description

Delete	Select this option to delete an existing IP interface.
VLAN	The VLAN associated with the IP interface. Only ports in this VLAN will be able to access the IP interface. This field is only available for input when creating a new interface.
IPv4 DHCP Enabled	Enable the DHCPv4 client by checking this box. If this option is enabled, the system will configure the IPv4 address and mask of the interface using the DHCPv4 protocol.
IPv4 DHCP Client Identifier Type	This specified which of the three types below, i.e. IfMac, ASCII or HEX, shall be used for the Client Identifier.
IPv4 DHCP Client Identifier IfMac	The interface name of DHCP client identifier. When DHCPv4 client is enabled and the client identifier type is 'ifmac', the configured interface's hardware MAC address will be used in the DHCP option 61 field.
IPv4 DHCP Client Identifier ASCII	The ASCII string of DHCP client identifier. When DHCPv4 client is enabled and the client identifier type is 'ascii', the ASCII string will be used in the DHCP option 61 field.
IPv4 DHCP Client Identifier HEX	The hexadecimal string of DHCP client identifier. When DHCPv4 client is enabled and the client identifier type 'hex', the hexadecimal value will be used in the DHCP option 61 field.
IPv4 DHCP Hostname	The hostname of DHCP client. If DHCPv4 client is enabled, the configured hostname will be used in the DHCP option 12 field. When this value is empty string, the field use the configured system name plus the latest three bytes of system MAC addresses as the hostname.
IPv4 DHCP Fallback Timeout	The number of seconds for trying to obtain a DHCP lease. After this period expires, a configured IPv4 address will be used as IPv4 interface address. A value of zero disables the fallback mechanism, such that DHCP will keep retrying until a valid lease is obtained. When fallback is disabled, no static IP address shall be specified. Legal values are 0 to 4294967295 seconds.
IPv4 DHCP Current Lease	For DHCP interfaces with an active lease, this column shows the current interface address, as provided by the DHCP server.
IPv4 Address	The IPv4 address of the interface in dotted decimal notation. If **DHCP** is enabled, this field configures the fallback address. The field may be left blank if IPv4 operation on the interface is not desired - or no DHCP fallback address is desired.
IPv4 Mask	The IPv4 network mask, in number of bits (prefix length). Valid values are **between 0 and 30 bits** for a IPv4 address. If **DHCP** is enabled, this field configures the fallback address network mask. The field may be left blank if IPv4 operation on the interface is not desired - or no DHCP fallback address is desired.
DHCPv6 Enable	Enable the DHCPv6 client by checking this box. If this option is enabled, the system will configure the IPv6 address of the interface using the DHCPv6 protocol.
DHCPv6 Rapid Commit	Enable the DHCPv6 Rapid-Commit option by checking this box. If this option is enabled, the DHCPv6 client terminates the waiting process as soon as a Reply message with a Rapid Commit option is received. This option is only manageable when DHCPv6 client is enabled.
DHCPv6 Current Lease	For DHCPv6 interface with an active lease, this column shows the interface address provided by the DHCPv6 server.
IPv6 Address	The IPv6 address of the interface. A IPv6 address is in 128-bit records represented as eight fields of up to four hexadecimal digits with a colon separating each field (:). For example, **fe80::215:c5ff:fe03:4dc7**. The symbol **::** is a special syntax that can be used as a shorthand way of representing multiple 16-bit groups of contiguous zeros; but it can appear only once. The system accepts the valid IPv6 unicast address only, except IPv4-Compatible address and IPv4-Mapped address. The field may be left blank if IPv6 operation on the interface is not desired.
IPv6 Mask	The IPv6 network mask, in number of bits (prefix length). Valid values are **between 1 and 128 bits** for a IPv6 address. The field may be left blank if IPv6 operation on the interface is not desired.
Resolving IPv6 DAD	The link-local address is formed from an interface identifier based on the hardware address which is supposed to be uniquely assigned. Once the DAD (Duplicate Address Detection) detects the address duplication, the operation on the interface SHOULD be disabled. At this moment, manual intervention is required to resolve the address duplication. For example, check whether the loop occurs in the VLAN or there is indeed other device occupying the same hardware address as the device in the VLAN. After making sure the specific link-local address is unique on the IPv6 link in use, delete and then add the specific IPv6 interface to restart the IPv6 operations on this interface.

The *IP Routes* section displays and manages the routing table, including destination networks, gateways, next hops, and route metrics. A maximum of 126 routes is supported.

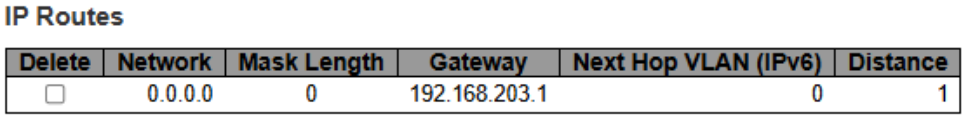

Parameter	Description
Delete	Select this option to delete an existing IP route.
Network	The destination IP network or host address of this route. Valid format is dotted decimal notation. A default route can use the value **0.0.0.0**.
Mask Length	The destination IP network or host mask, in number of bits (prefix length). It defines how much of a network address that must match, in order to qualify for this route. Valid values are **between 0 and 32 bits**. Only a default route will have a mask length of **0** (as it will match anything).
Gateway	The IP address of the IP gateway. Valid format is dotted decimal notation. In order to create a blackhole route (a route where all traffic to the network is discarded), use 'blackhole' verbatim.
Next Hop VLAN (Only for IPv6)	The VLAN ID (VID) of the specific IPv6 interface associated with the gateway. The given VID ranges from 1 to 4095 and will be effective only when the corresponding IPv6 interface is valid. If the IPv6 gateway address is link-local, it must specify the next hop VLAN for the gateway. If the IPv6 gateway address is not link-local, system ignores the next hop VLAN for the gateway.
Distance	The distance value of the route entry is used to provide the priority information of the routing protocols to routers. When two or more different routing protocols are involved and have the same destination, the distance value can be used to select the best path.

4.1.2 SNMP

Simple Network Management Protocol (SNMP) is a standardized protocol used to monitor and manage network devices. It enables a management system to read status information, receive event notifications, and modify configuration parameters on devices such as Ethernet switches using a structured Management Information Base (MIB).

4.1.2.1 System

This page configures the SNMP status.

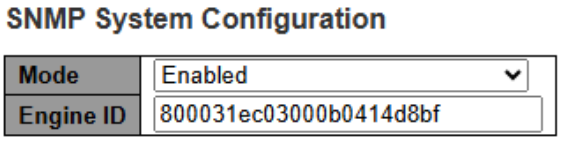

Parameter	Description
Mode	Indicates the SNMP mode operation. Possible modes are: **Enabled**: Enable SNMP mode operation. **Disabled**: Disable SNMP mode operation.
Engine ID	Indicates the SNMPv3 engine ID. The string must contain an even number (in hexadecimal format) with number of digits between 10 and 64, but all-zeros and all-'F's are not allowed. Only users on this Engine ID can access the device (local users), so changing the Engine ID will revoke access for all current local users.

4.1.2.2 Trap

SNMP traps are asynchronous notifications sent by a device to an SNMP manager to report events without polling. They are used to signal conditions such as link state changes, configuration errors, or fault conditions,

enabling timely network monitoring and response. SNMP Trap receivers can be added to be a trap destination on this page. Any SNMP server or management software that is configured to receive traps can be a SNMP Trap receiver.

Trap Destinations

A SNMP Trap Destination is the device that the SNMP Trap is sent to when it is triggered by an event. This is commonly a SNMP Server or a Network Management System.

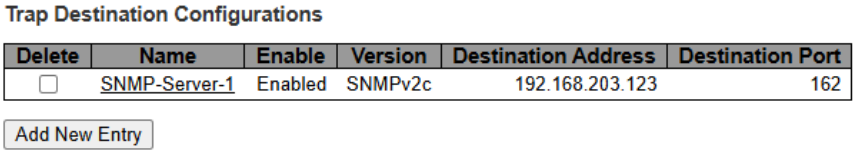

Parameter	Description
Name	Indicates the trap Configuration's name. Indicates the trap destination's name.
Enable	Indicates the SNMP trap mode operation. Possible modes are: **Enabled**: Enable SNMP trap mode operation. **Disabled**: Disable SNMP trap mode operation.
Version	Indicates the SNMP trap supported version. Possible versions are: **SNMPv1**: Set SNMP trap supported version 1. **SNMPv2c**: Set SNMP trap supported version 2c. **SNMPv3**: Set SNMP trap supported version 3.
Destination Address	Indicates the SNMP trap destination address. It allows a valid IP address in dotted decimal notation ('x.y.z.w'). And it also allows a valid hostname. A valid hostname is a string drawn from the alphabet (A-Za-z), digits (0-9), dot (.), dash (-). Spaces are not allowed, the first character must be an alpha character, and the first and last characters must not be a dot or a dash. Indicates the SNMP trap destination IPv6 address. IPv6 address is in 128-bit records represented as eight fields of up to four hexadecimal digits with a colon separating each field (:). For example, **fe80::215:c5ff:fe03:4dc7**. The symbol '**::**' is a special syntax that can be used as a shorthand way of representing multiple 16-bit groups of contiguous zeros; but it can appear only once. It can also represent a legally valid IPv4 address. For example, '**::192.1.2.34**'.
Destination Port	Indicates the SNMP trap destination port. SNMP Agent will send SNMP message via this port, the port range is **1~65535**.

Trap Sources

The trap sources which trigger a SNMP trap can be set by OID. A trap is sent for the given trap source if at least one filter with filter type included matches the filter, and no filters with filter type excluded matches. The maximum entry count is 32.

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
Name	Indicates the name for the entry.
Type	The filter type for the entry. Possible types are: **included**: An optional flag to indicate a trap is sent for the given trap source is matched. **excluded**: An optional flag to indicate a trap is not sent for the given trap source is matched.

Subset OID	The IP address of the IP gateway. Valid format is dotted decimal notation. In order to create a blackhole route (a route where all traffic to the network is discarded), use 'blackhole' verbatim.

4.1.2.3 Communities

SNMP communities are shared text strings used for access control between an SNMP manager and a device. They define read-only or read-write permissions and must match on both sides for SNMP requests to be accepted.

SNMPv3 Community Configuration

Delete	Community name	Community secret	Source IP	Source Prefix
☐	public	secret	192.168.203.0	24

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
Community Name	Indicates the security name to map the community to the SNMP Groups configuration. The allowed string length is 1 to 32, and the allowed content is ASCII characters from 33 to 126.
Community Secret	Indicates the community secret (access string) to permit access using SNMPv1 and SNMPv2c to the SNMP agent. The allowed string length is 1 to 32, and the allowed content is ASCII characters from 33 to 126.
Source IP	Indicates the SNMP access source address. A particular range of source addresses can be used to restrict source subnet when combined with source prefix.
Source Prefix	Indicates the SNMP access source address prefix. This is the subnet mask in CIDR notation.

4.1.2.4 Users

SNMPv3 users define authenticated and optionally encrypted access to a device. Each user is associated with authentication (MD5 or SHA) and privacy (DES or AES) settings, providing secure management and access control.

SNMPv3 User Configuration

Delete	Engine ID	User Name	Security Level	Authentication Protocol	Authentication Password	Privacy Protocol	Privacy Password
☐	800031ec03000b0414d8bf	User1	Auth, Priv	MD5		DES	

Add New Entry | Save | Reset

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
Engine ID	An octet string identifying the engine ID that this entry should belong to. The string must contain an even number(in hexadecimal format) with number of digits between 10 and 64, but all-zeros and all-'F's are not allowed. The SNMPv3 architecture uses the User-based Security Model (USM) for message security and the View-based Access Control Model (VACM) for access control. For the USM entry, the usmUserEngineID and usmUserName are the entry's keys. In a simple agent, usmUserEngineID is always that agent's own snmpEngineID value. The value can also take the value of the snmpEngineID of a remote SNMP engine with which this user can communicate. In other words, if user engine ID equal system engine ID then it is local user; otherwise it's remote user.
User Name	A string identifying the user name that this entry should belong to. The allowed string length is 1 to 32, and the allowed content is ASCII characters from 33 to 126.

Security Level	Indicates the security model that this entry should belong to. Possible security models are: **NoAuth, NoPriv**: No authentication and no privacy. **Auth, NoPriv**: Authentication and no privacy. **Auth, Priv**: Authentication and privacy. The value of security level cannot be modified if entry already exists. That means it must first be ensured that the value is set correctly.
Authentication Protocol	Indicates the authentication protocol that this entry should belong to. Possible authentication protocols are: **None**: No authentication protocol. **MD5**: An optional flag to indicate that this user uses MD5 authentication protocol. **SHA**: An optional flag to indicate that this user uses SHA authentication protocol. The value of security level cannot be modified if entry already exists. That means must first ensure that the value is set correctly.
Authentication Password	A string identifying the authentication password phrase. For MD5 authentication protocol, the allowed string length is 8 to 32. For SHA authentication protocol, the allowed string length is 8 to 40. The allowed content is ASCII characters from 33 to 126.
Privacy Protocol	Indicates the privacy protocol that this entry should belong to. Possible privacy protocols are: **None**: No privacy protocol. **DES**: An optional flag to indicate that this user uses DES authentication protocol. **AES**: When available, an optional flag to indicate that this user uses AES authentication protocol.
Privacy Password	A string identifying the privacy password phrase. The allowed string length is 8 to 32, and the allowed content is ASCII characters from 33 to 126.

4.1.2.5 Groups

This defines a user group which can then collectively be assigned access rights for SNMP operations.

SNMPv3 Group Configuration

Delete	Security Model	Security Name	Group Name
☐	v1	public	default_ro_group
☐	v1	private	default_rw_group
☐	v2c	public	default_ro_group
☐	v2c	private	default_rw_group

Add New Entry | Save | Reset

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
Security Model	Indicates the security model that this entry should belong to. Possible security models are: **v1**: Reserved for SNMPv1. **v2c**: Reserved for SNMPv2c. **usm**: SNMPv3, User-based Security Model (USM).
User Name	A string identifying the user name that this entry should belong to. The allowed string length is 1 to 32, and the allowed content is ASCII characters from 33 to 126.
Security Name	A string identifying the security name that this entry should belong to. The allowed string length is 1 to 32, and the allowed content is ASCII characters from 33 to 126.
Group Name	A string identifying the group name that this entry should belong to. The allowed string length is 1 to 32, and the allowed content is ASCII characters from 33 to 126.

4.1.2.6 Views

SNMP views define subsets of MIB objects that can be accessed by specific groups. They allow filtering which objects are visible for read or write operations, supporting fine-grained access control.

SNMPv3 View Configuration

Delete	View Name	View Type	OID Subtree
☐	default_view	included	.1

Add New Entry | Save | Reset

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
View Name	A string identifying the view name that this entry should belong to. The allowed string length is 1 to 32, and the allowed content is ASCII characters from 33 to 126.
View Type	Indicates the view type that this entry should belong to. Possible view types are: **included**: An optional flag to indicate that this view subtree should be included. **excluded**: An optional flag to indicate that this view subtree should be excluded. In general, if a view entry's view type is 'excluded', there should be another view entry existing with view type as 'included' and it's OID subtree should overstep the 'excluded' view entry.
OID Subtree	A string identifying the security name that this entry should belong to. The allowed string length is 1 to 32, and the allowed content is ASCII characters from 33 to 126.
Group Name	The OID defining the root of the subtree to add to the named view. The allowed OID length is 1 to 128. The allowed string content is digital number or asterisk (*).

4.1.2.7 Access

The SNMPv3 access table specifies which users or groups can perform read or write operations on defined SNMP views. It enforces security and permissions for managing MIB objects.

SNMPv3 Access Configuration

Delete	Group Name	Security Model	Security Level	Read View Name	Write View Name
☐	default_ro_group	any	NoAuth, NoPriv	default_view	None
☐	default_rw_group	any	NoAuth, NoPriv	default_view	default_view

Add New Entry | Save | Reset

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
Group Name	A string identifying the group name that this entry should belong to. The allowed string length is 1 to 32, and the allowed content is ASCII characters from 33 to 126.
Security Model	Indicates the security model that this entry should belong to. Possible security models are: **any**: Any security model accepted(v1|v2c|usm). **v1**: Reserved for SNMPv1. **v2c**: Reserved for SNMPv2c. **usm**: SNMPv3, User-based Security Model (USM).
Security Level	Indicates the security model that this entry should belong to. Possible security models are: **NoAuth, NoPriv**: No authentication and no privacy. **Auth, NoPriv**: Authentication and no privacy. **Auth, Priv**: Authentication and privacy.
Read View Name	The name of the MIB view defining the MIB objects for which this request may request the current values. The allowed string length is 1 to 32, and the allowed content is ASCII characters from 33 to 126.
Write View Name	The name of the MIB view defining the MIB objects for which this request may potentially set new values. The allowed string length is 1 to 32, and the allowed content is ASCII characters from 33 to 126.

4.1.3 NTP

NTP synchronizes the system clock of network devices with reference time sources, ensuring consistent timestamps for logs, events, and coordinated operations across the network.

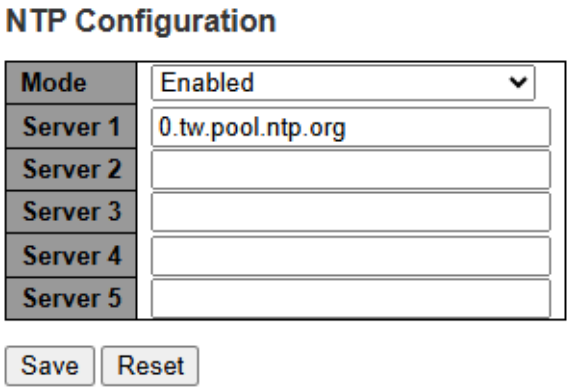

Parameter	Description
Mode	Indicates the NTP mode operation. Possible modes are: **Enabled**: Enable NTP client mode operation. **Disabled**: Disable NTP client mode operation.
Server #	Provide the IPv4 or IPv6 address of a NTP server. IPv6 address is in 128-bit records represented as eight fields of up to four hexadecimal digits with a colon separating each field (:). For example, 'fe80::215:c5ff:fe03:4dc7'. The symbol '::' is a special syntax that can be used as a shorthand way of representing multiple 16-bit groups of contiguous zeros; but it can appear only once. It can also represent a legally valid IPv4 address. For example, '::192.1.2.34'. In addition, it can also accept a domain name address.

4.1.4 Time

Time Zone

Configure the time zone to be used for system time and system logs.

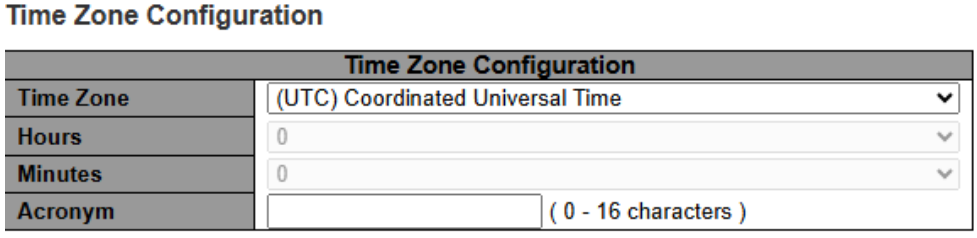

Parameter	Description
Time Zone	Lists various Time Zones worldwide. Select appropriate Time Zone from the drop down and click Save to set. The *Manual Setting* options is used for the specific time zone which is excluded from the options list.
Hours	Number of hours offset from UTC. The field only available when time zone manual setting.
Minutes	Number of minutes offset from UTC. The field only available when time zone manual setting.
Acronym	User can set the acronym of the time zone. This is a User configurable acronym to identify the time zone. (Range: Up to 16 characters) Notice the string `''` is a special syntax that is reserved for null input.

Daylight Saving Time

This is used to set the clock forward or backward according to the configurations set below for a defined Daylight-Saving Time duration. Select **`Disable`** to disable the Daylight-Saving Time configuration. Select **`Recurring`** and configure the Daylight-Saving Time duration to repeat the configuration every year. Select **`Non-Recurring`** and configure the Daylight-Saving Time duration for single time configuration. (Default: **`Disabled`**)

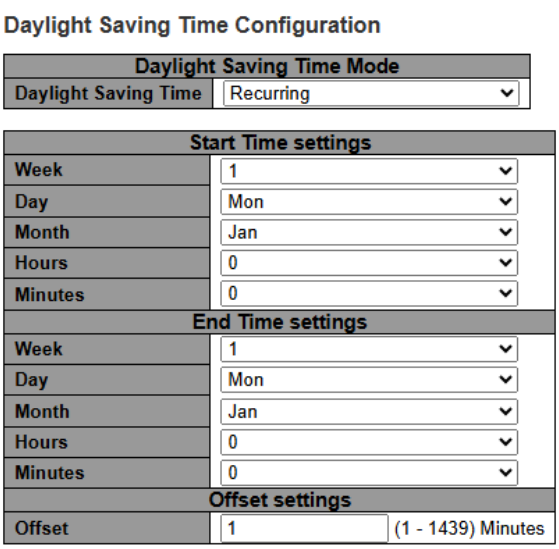

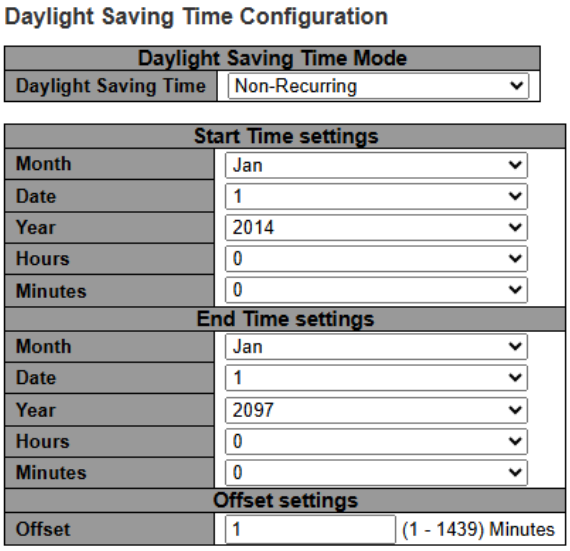

Parameter	Description
Start Time settings	**`Recurring`**: Select the starting week, day, month, hour, and minutes of the daylight-saving time. **`Non-Recurring`**: Select the starting month, day, year, hour, and minutes of the daylight-saving time.
End Time settings	**`Recurring`**: Select the ending week, day, month, hour, and minutes of the daylight-saving time. **`Non-Recurring`**: Select the ending month, day, year, hour, and minutes of the daylight-saving time.
Offset	Enter the number of minutes to add during Daylight Saving Time. (Range: **`1 to 1439`**)

4.1.5 Log

Configure the sensitivity of the system log as well as sending syslog messages to a server.
Logs can be viewed under Monitor→System→Log.

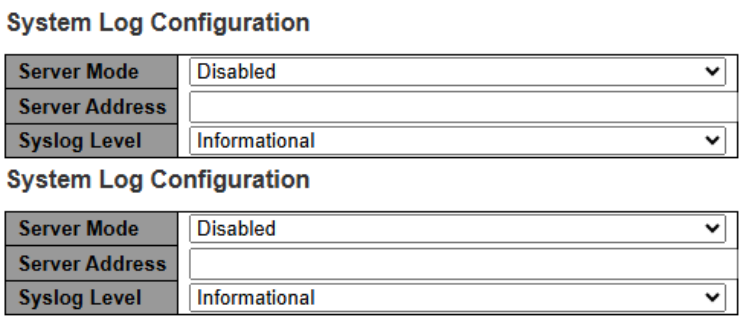

Parameter	Description
Server Mode	Indicates the server mode operation. When the mode operation is enabled, the syslog message will send out to syslog server. The syslog protocol is based on UDP communication and received on UDP

	port 514 and the syslog server will not send acknowledgments back sender since UDP is a connectionless protocol and it does not provide acknowledgments. The syslog packet will always send out even if the syslog server does not exist. Possible modes are: **Enabled**: Enable server mode operation. **Disabled**: Disable server mode operation.
Server Address	Indicates the host address of syslog server.
Syslog Level	Indicates what kind of message will send to syslog server. Possible modes are: **Error**: Send specific messages which severity code is less or equal than Error (3). **Warning**: Send specific messages which severity code is less or equal than Warning (4). **Notice**: Send specific messages which severity code is less or equal than Notice (5). **Informational**: Send specific messages which severity code is less or equal than Informational (6).

4.2 Green Ethernet (Port Power Savings)

This feature controls settings for Energy Efficient Ethernet (EEE), which is defined in IEEE 802.3az, and additional power saving features. EEE is a power saving option that reduces the power usage when there is low or no traffic utilization. A detailed view of the Port Power Saving status is available at Monitor→Green Ethernet→Port Power Savings.

EEE works by powering down circuits when there is no traffic. When a port gets data to be transmitted, all circuits are powered up. The time it takes to power up the circuits is named wakeup time. The default wakeup time is 17 us for 1Gbit links and 30 us for other link speeds. EEE devices must agree upon the value of the wakeup time to make sure that both the receiving and transmitting device has all circuits powered up when traffic is transmitted. The devices can exchange wakeup time information using the LLDP protocol.

EEE works for ports in auto-negotiation mode, where the port is negotiated to either 1G or 100 Mbit in full duplex mode. For ports that are not EEE-capable, the corresponding EEE checkboxes are grayed out and thus impossible to enable EEE for.

When a port is powered down for saving power, outgoing traffic is stored in a buffer until the port is powered up again. Because there are some overhead in turning the port down and up, more power can be saved if the traffic can be buffered up until a large burst of traffic can be transmitted. Buffering traffic will introduce some latency in the traffic.

The switch can be set to optimize EEE for either best power saving or least traffic latency.

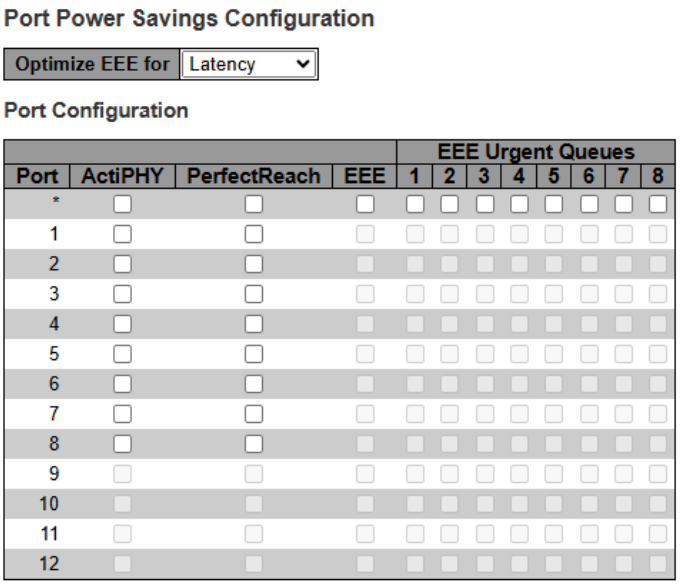

Parameter	Description
Port	The switch port number of the logical port.
ActiPHY	Link down power savings enabled. ActiPHY works by lowering the power for a port when there is no link. The port is power up for short moment in order to determine if cable is inserted.
PerfectReach	Cable length power savings enabled. PerfectReach works by determining the cable length and lowering the power for ports with short cables.
EEE	Controls whether EEE is enabled for this switch port. For maximizing power savings, the circuit isn't started at once transmit data is ready for a port, but is instead queued until a burst of data is ready to be transmitted. This will give some traffic latency. If desired it is possible to minimize the latency for specific frames, by mapping the frames to a specific queue (done with QoS), and then mark the queue as an urgent queue. When an urgent queue gets data to be transmitted, the circuits will be powered up at once and the latency will be reduced to the wakeup time.
EEE Urgent Queues	Queues set will activate transmission of frames as soon as data is available. Otherwise, the queue will postpone transmission until a burst of frames can be transmitted.

4.3 Thermal Protection

Thermal protection is used to protect the hardware from getting overheated. The temperature and port protection status can be viewed at Monitor→Thermal Protection.

Temperature settings for groups Port groups

Group	Temperature
0	255 °C
1	255 °C
2	255 °C
3	255 °C

Port	Group
*	<>
1	Disabled
2	Disabled
3	Disabled
4	Disabled
5	Disabled
6	Disabled
7	Disabled
8	Disabled
9	Disabled
10	Disabled
11	Disabled
12	Disabled

When the temperature exceeds the configured thermal protection temperature, ports will be turned off in order to decrease the power consumption. It is possible to arrange the ports with different groups. Each group can be given a temperature at which the corresponding ports shall be turned off.

The temperature at which a port will be turned off can be defined for 4 separate groups. The ports can then be assigned to one of the groups to be shut down when they reach the corresponding temperature threshold. Temperatures between 0 and 255 C are supported.

4.4 Ports

The port configuration contains settings for the RJ45 and SFP Ethernet ports. Detailed port information can be viewed in the Monitor→Ports section.

Port Configuration — Refresh

Port	Link	Warning	Speed		Adv Duplex		Adv speed						Flow Control			PFC		Maximum Frame Size	Excessive Collision Mode	Frame Length Check	FEC Mode
			Current	Configured	Fdx	Hdx	10M	100M	1G	2.5G	5G	10G	Enable	Curr Rx	Curr Tx	Enable	Priority				
*				<>	☑	☑	☑	☑	☑	☐	☐	☐	☐			☐	0-7	10240	<>	☐	<>
1			Down	Automatic	☑	☑	☑	☑	☑	☐	☐	☐	☐	✕	✕	☐	0-7	10240	Discard	☐	
2			Down	Automatic	☑	☑	☑	☑	☑	☐	☐	☐	☐	✕	✕	☐	0-7	10240	Discard	☐	
3			Down	Automatic	☑	☑	☑	☑	☑	☐	☐	☐	☐	✕	✕	☐	0-7	10240	Discard	☐	
4			Down	Automatic	☑	☑	☑	☑	☑	☐	☐	☐	☐	✕	✕	☐	0-7	10240	Discard	☐	
5			Down	Automatic	☑	☑	☑	☑	☑	☐	☐	☐	☐	✕	✕	☐	0-7	10240	Discard	☐	
6			Down	Automatic	☑	☑	☑	☑	☑	☐	☐	☐	☐	✕	✕	☐	0-7	10240	Discard	☐	
7			Down	Automatic	☑	☑	☑	☑	☑	☐	☐	☐	☐	✕	✕	☐	0-7	10240	Discard	☐	
8			1Gfdx	Automatic	☑	☑	☑	☑	☑	☐	☐	☐	☐	✕	✕	☐	0-7	10240	Discard	☐	
9			Down	Automatic	☑	☐	☐	☑	☑	☑	☑	☑	☐	✕	✕	☐	0-7	10240		☐	aut
10			Down	Automatic	☑	☐	☐	☑	☑	☑	☑	☑	☐	✕	✕	☐	0-7	10240		☐	aut
11			Down	Automatic	☑	☐	☐	☑	☑	☑	☑	☑	☐	✕	✕	☐	0-7	10240		☐	aut
12			Down	Automatic	☑	☐	☐	☑	☑	☑	☑	☑	☐	✕	✕	☐	0-7	10240		☐	aut

Parameter	Description
Port	This is the logical port number for this row.
Link	The current link state is displayed graphically. Green: Link up Red: Link down
Warning	Operational warnings of the port. Gray: No warnings Yellow: There are warnings, use tooltip by hovering over the icon to see.
Current Link Speed	Provides the current link speed of the port.
Configured Link Speed	Selects any available link speed for the given switch port. Only speeds supported by the specific port is shown. Possible speeds are:

	Disabled - Disables the switch port operation. **Automatic** - Port auto negotiating speed and duplex with the link partner and selects the highest speed that is compatible with the link partner. **10Mbps HDX** - Forces the port in 10Mbps half duplex mode. **10Mbps FDX** - Forces the port in 10Mbps full duplex mode. **100Mbps HDX** - Forces the port in 100Mbps half duplex mode. **100Mbps FDX** - Forces the port in 100Mbps full duplex mode. **1Gbps FDX** - Forces the port in 1Gbps full duplex **2.5Gbps FDX** - Forces the port in 2.5Gbps full duplex mode. **5Gbps FDX** - Forces the port in 5Gbps full duplex mode. **10Gbps FDX** - Forces the port in 10Gbps full duplex mode.
Advertise Duplex	When duplex is set as auto i.e auto negotiation, the port will only advertise the specified duplex as either **Fdx** or **Hdx** to the link partner. By default, port will advertise all the supported duplexes if the Duplex is Auto. This setting is only available for pure copper ports.
Advertise Speed	When Speed is set as auto i.e auto negotiation, the port will only advertise the specified speeds (**10M 100M 1G 2.5G 5G 10G**) to the link partner. By default, port will advertise all the supported speeds if speed is set as **Automatic**. This setting is only available for pure copper ports.
Flow Control	When Auto Speed is selected on a port, this section indicates the flow control capability that is advertised to the link partner. When a fixed-speed setting is selected, that is what is used. The Current Rx column indicates whether pause frames on the port are obeyed, and the Current Tx column indicates whether pause frames on the port are transmitted. The Rx and Tx settings are determined by the result of the last Auto Negotiation. Check the configured column to use flow control. This setting is related to the setting for Configured Link Speed. NOTICE: The 100FX standard does not support Auto Negotiation, so when in 100FX mode the flow control capabilities will always be shown as **disabled**.
PFC	When PFC (802.1Qbb Priority Flow Control) is enabled on a port then flow control on a priority level is enabled. Through the *Priority* field, range (one or more) of priorities can be configured, e.g. '0-3,7' which equals '0,1,2,3,7'. PFC is not supported through auto negotiation. PFC and Flow Control cannot both be enabled on the same port.
Maximum Frame Size	Enter the maximum frame size allowed for the switch port, including FCS. The range is **1518-10240** bytes.
Excessive Collision Mode	Configure port transmit collision behavior. **Discard:** Discard frame after 16 collisions (default). **Restart:** Restart backoff algorithm after 16 collisions.
Frame Length Check	Configures if frames with incorrect frame length in the EtherType/Length field shall be dropped. An Ethernet frame contains a field EtherType which can be used to indicate the frame payload size (in bytes) for values of 1535 and below. If the EtherType/Length field is above 1535, it indicates that the field is used as an EtherType (indicating which protocol is encapsulated in the payload of the frame). If "frame length check" is enabled, frames with payload size less than 1536 bytes are dropped, if the EtherType/Length field does not match the actually payload length. If "frame length check" is disabled, frames are not dropped due to frame length mismatch. Note: No drop counters count frames dropped due to frame length mismatch
FEC	FEC is short for Forward Error Correction. It is a technique for controlling errors over an unreliable link. The idea is that the sender adds some extra bits to the frame that allows a receiver to correct bit errors in the received frame. R-FEC (IEEE802.3 clause 74 - sometimes called Firecode). This is meant for 10G. The parameter affects both what is requested during clause 73 aneg and what the port is configured to use if not running clause 73 aneg. If running clause 73 aneg on 10G ports we always tell the link partner that we support R-FEC. What the end user can control with the fec command is whether we request R-FEC. If either us or the link partner requests R-FEC, the port will end up using R-FEC. **auto:** This is the default and means the following: If a 10G port runs clause 73, R-FEC will be requested. Otherwise, no FEC will be enabled. **r-fec:** If a 10G port runs clause 73, only R-FEC will be requested. If a 10G port does not run clause 73, but is loaded with at least a 10G SFP and the speed is at least 5G, only R-FEC will be enabled. Otherwise, no FEC will be enabled.

	none : If the port is running clause 73, R-FEC will not be requested (but remember that this does not mean that the clause 73 aneg will not result in the port running FEC). Otherwise, the port will not run any FEC.

4.5 Digital Input/Output

The 9561 provides a set of industrial-control contacts for connecting external alarm devices. DI1 and DI2 are digital inputs that read the open or closed state of an external dry-contact device. DO is a dry-contact relay output that gives a local, immediate response to a DI event. ALARM is a separate dry-contact relay output that aggregates system alarms, such as a power or redundant-power loss, for uplink reporting.

This section is organised in three layers: the operating concept, the configuration workflow from wiring through CLI, and a set of practical use cases that map DI/DO to common on-site needs.

Contact	Type	Direction	Purpose
DI1	Digital input	External device to switch	Reads the open or closed state of external contact 1.
DI2	Digital input	External device to switch	Reads the open or closed state of external contact 2.
DO	Digital output	Switch to external device	Drives an external relay or alarm device on a DI event.
ALARM	Alarm relay	Switch to external device	Independent relay driven automatically by any active eligible system alarm.

4.5.1 Operation Concept

A simple way to understand the DI/DO path is to keep four roles in mind:

DI is what the switch sees. DI1 and DI2 read the state changes of external contacts: a door opened, smoke detected, a cabinet forced open, and so on.

DO is the local immediate response. DO is driven by the result of the named alarm expression evaluated on DI1 or DI2. When a DI event makes its named alarm become Active, DO closes immediately to drive an on-site buzzer or warning light. The whole path is gated by the DI1/DI2 Diagnostic and Relay DO settings.

The named alarm expression is the monitoring rule. The alarm framework is the switch's generic alarm engine. It accepts any boolean alarm source (a DI status, a port-link status, a sensor flag, and so on) and binds one or more sources to a named alarm. The framework then continuously monitors the bound sources and updates the alarm's Active or Clear state. By convention the expression must evaluate to TRUE in the normal state; when it flips to FALSE the alarm becomes Active.

The ALARM relay is the uplink. The ALARM relay closes when any eligible system alarm is Active — for example power, redundant-power, voltage or temperature alarms. DI1 / DI2 alarms drive the DO relay, not the ALARM relay (Reboot and Factory-reset alarms do not drive it either). It is typically used to escalate such system alarms upstream.

Once a named alarm is Active it can be consumed in several ways: it is listed in show alarm status, it lights the front-panel ALARM LED, it can drive the ALARM relay (when Relay Alarm is enabled), and it can act as a source for SNMP-trap or syslog notifications (each enabled separately).

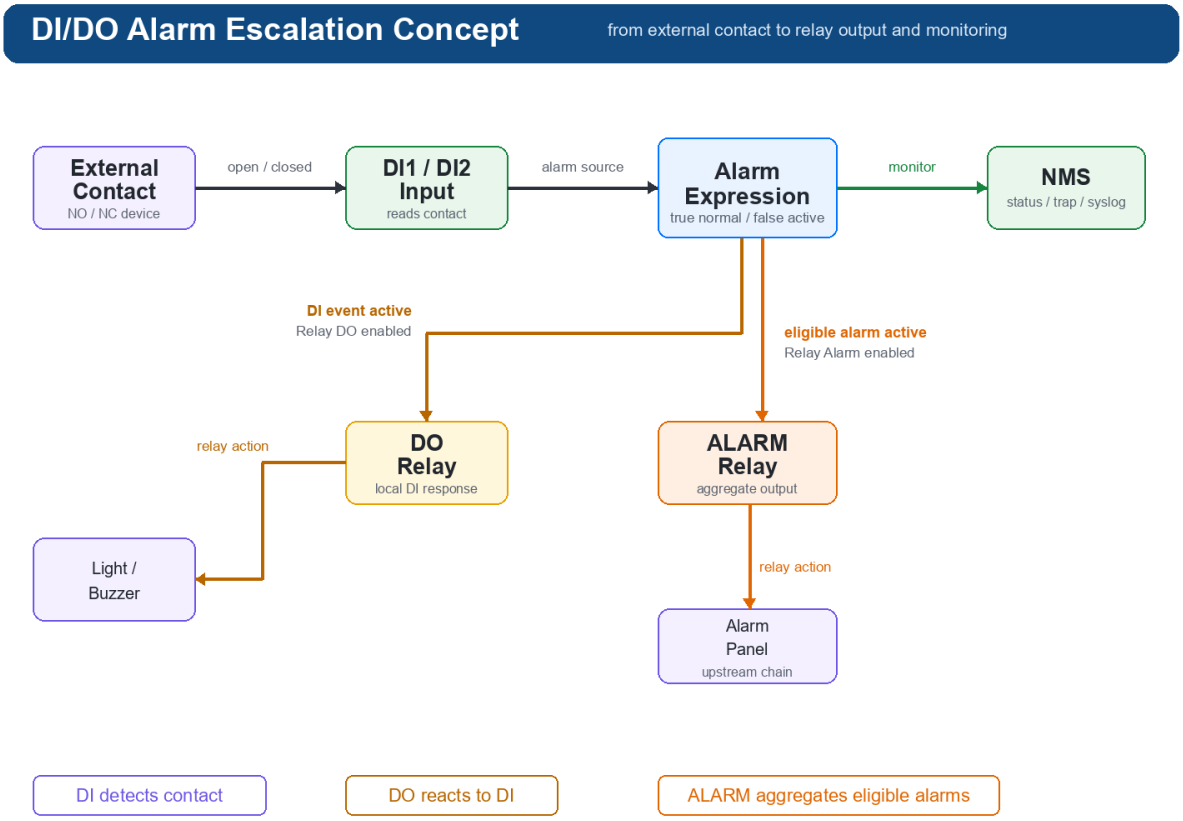

Figure: Conceptual block diagram of the DI / DO / ALARM signal path.

Operational note - escalation delay. From the moment an external contact changes state to the moment the relay closes or a notification is sent, expect roughly 4 to 5 seconds. The delay comes from DI debounce (the input must read the same abnormal state on five consecutive samples before the alarm is raised) plus the alarm framework's periodic scan. Allow for this delay when designing contact-closure event timing.

4.5.2 Relay Physical Location

The ALARM relay and the DO relay are separate physical relay outputs. The ALARM relay terminal block is located near the power/alarm terminal area. The DO relay terminal is part of the digital input/output terminal block.

Figure: Physical location and appearance of the ALARM relay terminal block.

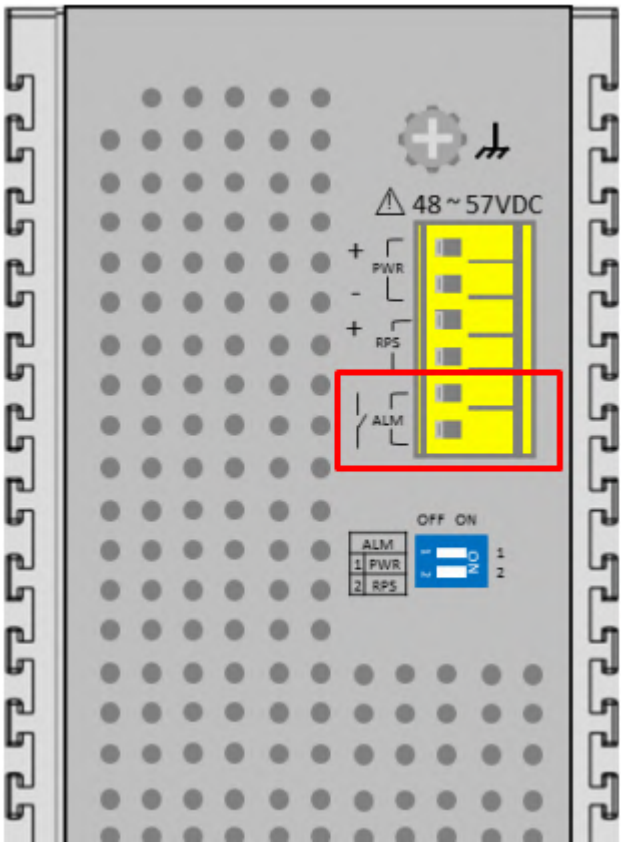

Figure: Physical location and appearance of the DO relay terminal block.

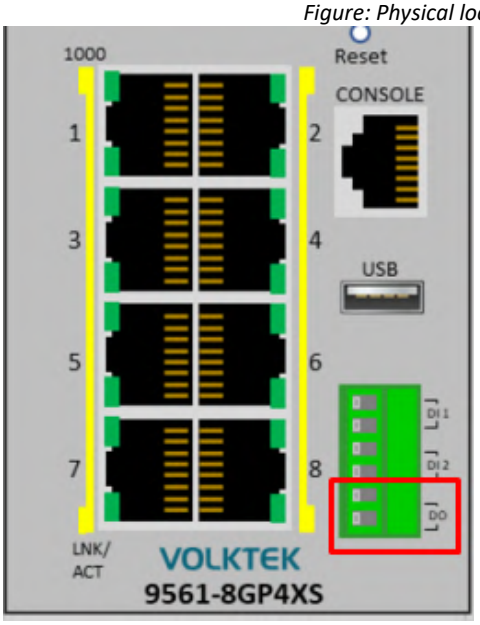

4.5.3 External Contact Type and DI Attach Mode

NO (Normal Open) and NC (Normal Close) describe the contact type of the external device itself, not a different wiring style. In both cases you connect the two leads of the external contact to the DI terminal pair; no separate ground wire is needed. What differs is how the device's internal contact behaves between the normal and event states. After wiring, the switch must be told which type is connected (the DI Attach Mode, or the di attachment setting), otherwise the alarm logic is inverted.

Contact type	Normal state	Abnormal state	Typical external devices

NO (Normal Open)	Contact open	Contact closed	Smoke detector, push-button, water-leak sensor, photoelectric sensor
NC (Normal Close)	Contact closed	Contact open	Emergency-stop button, door magnetic switch, cabinet tamper switch

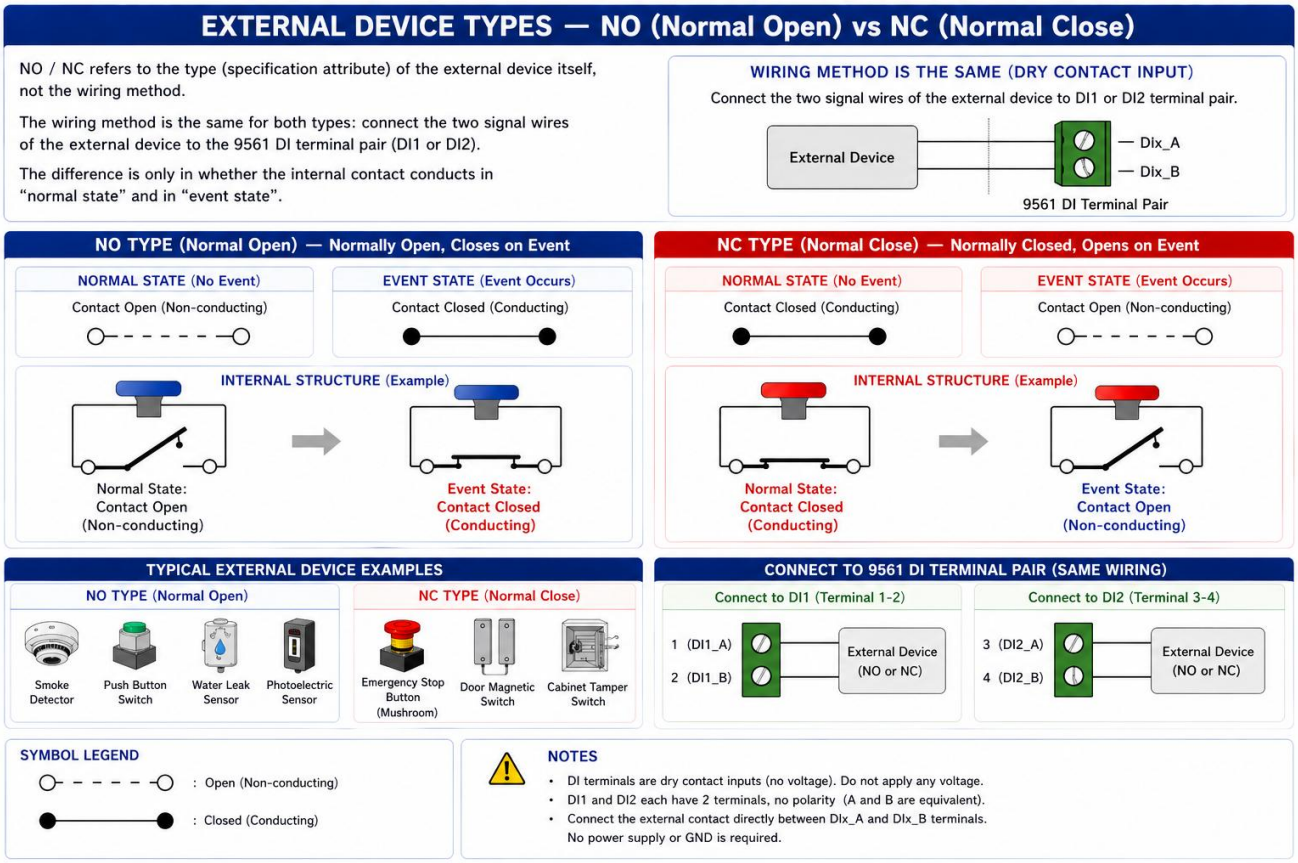

The raw DI value depends only on the physical state of the contact (1 = closed, 0 = open). The switch derives the trigger status by comparing the raw value against the configured attach mode:

Contact type	State	External contact	DI Value	Trigger status
NO	Normal	Open	0	Not triggered
NO	Abnormal	Closed	1	Triggered
NC	Normal	Closed	1	Not triggered
NC	Abnormal	Open	0	Triggered

Figure: External device types showing NO (Normal Open) and NC (Normal Close) behavior, same dry-contact wiring method, DI terminal pair connection, typical examples, and notes.

4.5.4 Configuration Workflow

The complete flow uses only three command families: di1/di2 attachment (or the Web DI Attach Mode) to declare the contact type; alarm <name> <expression> to create the monitoring rule; and relay enable do / relay enable alarm to enable the relay outputs. The diagram below shows the end-to-end sequence.

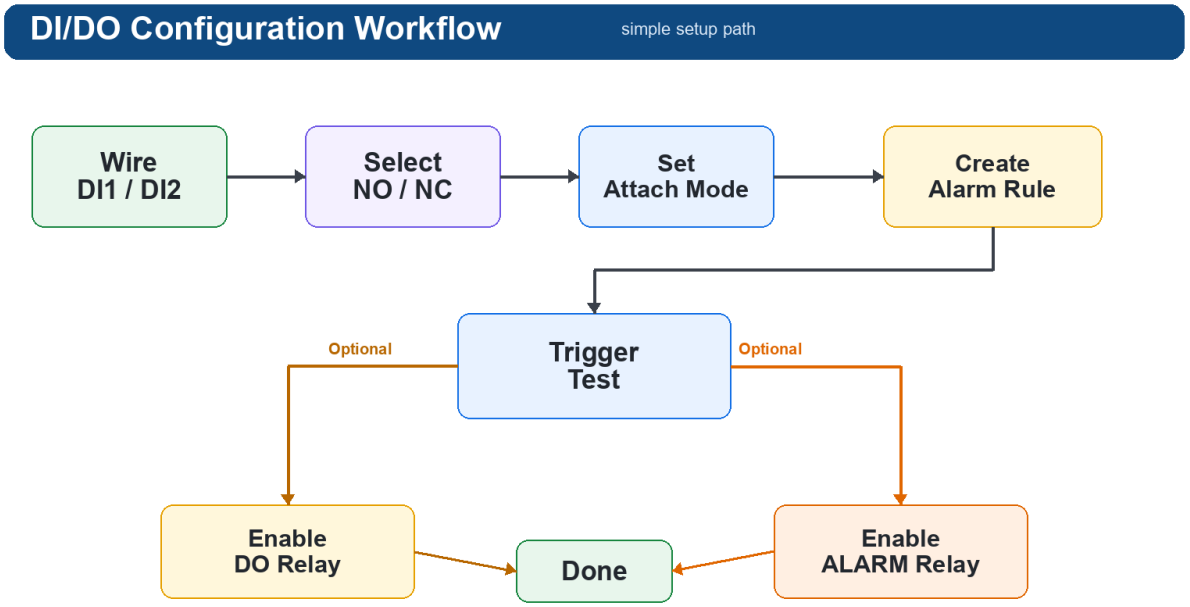

Figure: DI/DO configuration workflow, from wiring to enabling the relay outputs.

4.5.5 Web Operation

Open Configuration > DI/DO to set the relay gates and the DI diagnostic attach modes. Enable Relay Alarm for the ALARM relay path and Relay DO for the DO relay path, and set the DI1/DI2 Diagnostic attach mode (Normal Open or Normal Close) to match the connected device.

Figure: Configuration > DI/DO page showing the relay gates and DI diagnostic attach modes.

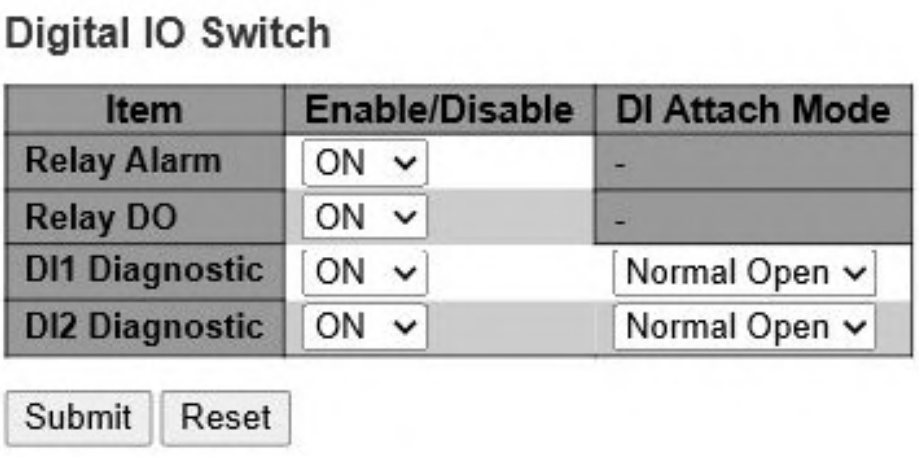

Parameter	Description
Item	DI/DO function row on the page, such as Relay Alarm, Relay DO, DI1 Diagnostic, or DI2 Diagnostic.
Enable/Disable	Enables or disables the selected relay gate or DI diagnostic function.

DI Attach Mode	Selects the external contact type for a DI diagnostic input. This setting is available for DI1 Diagnostic and DI2 Diagnostic only.
Relay Alarm	Enables the ALARM relay output path for eligible system alarms.
Relay DO	Enables the DO relay output path for DI events.
DI1 Diagnostic	Enables DI1 diagnostic monitoring and selects the DI1 attach mode.
DI2 Diagnostic	Enables DI2 diagnostic monitoring and selects the DI2 attach mode.
Submit	Applies the DI/DO configuration changes.
Reset	Restores the unsaved page values.

4.5.6 Common Use Cases

The four cases below cover the most common DI/DO patterns. Each shows the configuration and the resulting signal path.

4.5.6.1 Case 1 - Server-room door monitoring (NC)

A door magnetic switch (NC) is wired to DI1. When the door is opened unexpectedly, the switch reports the event to the NMS and drives a local warning light or buzzer through DO.

```
configure terminal
 di1 attachment normal-close
 alarm alarm.sasi.DI1 sasi.status@DI1==true
 relay enable do
end
```

Both the DI1 Diagnostic and Relay DO gates must be enabled. When the door opens, DI1 becomes triggered, the alarm.sasi.DI1 expression evaluates false, and the named alarm goes Active: DO closes to energise the on-site warning light, and the alarm appears in show alarm status as an SNMP-trap or syslog source. If only NMS reporting is needed, omit relay enable do.

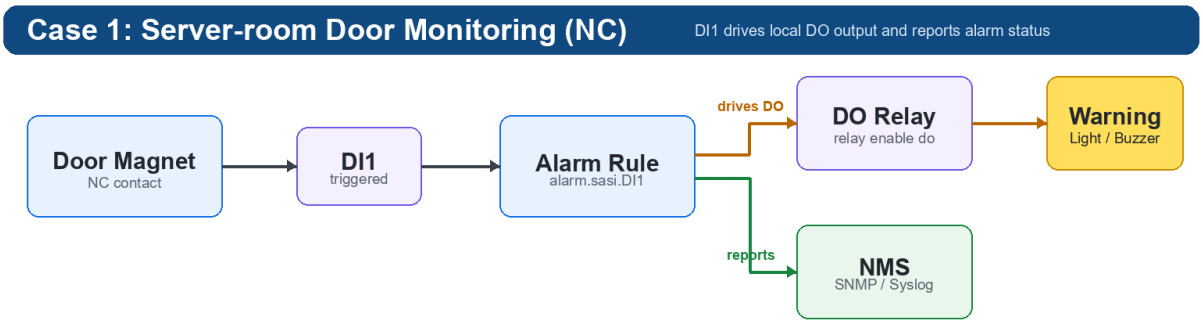

Figure: Case 1 - server-room door monitoring signal path.

4.5.6.2 Case 2 - Smoke detection with external buzzer (NO)

A smoke detector (NO) is wired to DI1. When smoke is detected, DO drives an on-site buzzer.

```
configure terminal
 di1 attachment normal-open
 alarm alarm.sasi.DI1 sasi.status@DI1==false
 relay enable do
end
```

All three settings are required: di1 attachment normal-open declares the contact type, the alarm command registers the named alarm alarm.sasi.DI1 (a monitored alarm source) over sasi.status@DI1, and relay enable do enables the physical DO closure. When smoke is present, the expression evaluates false, the alarm goes Active, and DO closes to sound the buzzer.

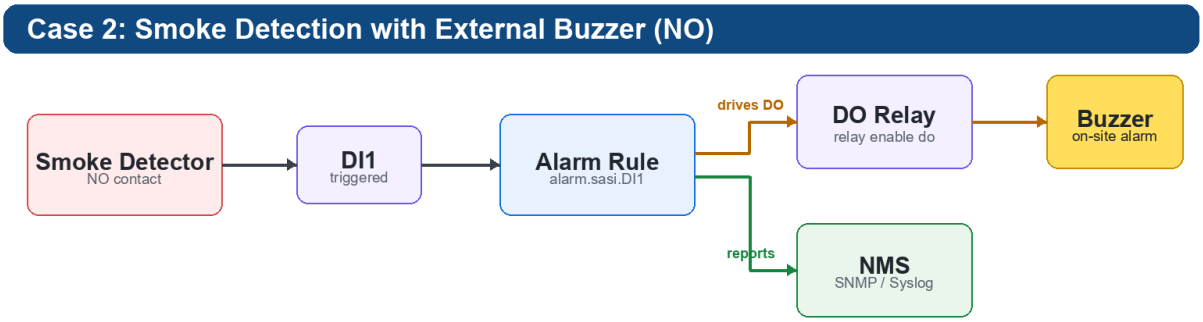

Figure: Case 2 - smoke detection with external buzzer signal path.

4.5.6.3 Case 3 - Redundant-power (RPS) loss with ALARM relay

The unit has a redundant power feed (RPS). If the RPS feed is lost, the switch reports it to the NMS and drives an on-site warning light / upstream relay through the ALARM relay. The ALARM relay is the correct output here because it carries system alarms (power, voltage, temperature) - DI1 / DI2 inputs do not drive it (they drive DO).

```
configure terminal
 alarm alarm.sasi.NO_RPS sasi.status@NO_RPS==false
 relay enable alarm
end
```

The on-board RPS Absent DIP must be ON for the NO_RPS diagnostic to be raised (a physical switch, no CLI/Web), and the Relay Alarm gate (relay enable alarm) must be enabled. The ALARM relay is driven by eligible system alarms (board temperature, voltage rails, NO_PWR, NO_RPS); DI1 / DI2, Reboot and Factory-reset are excluded - DI1 / DI2 drive the DO relay. If only NMS reporting is needed, omit relay enable alarm.

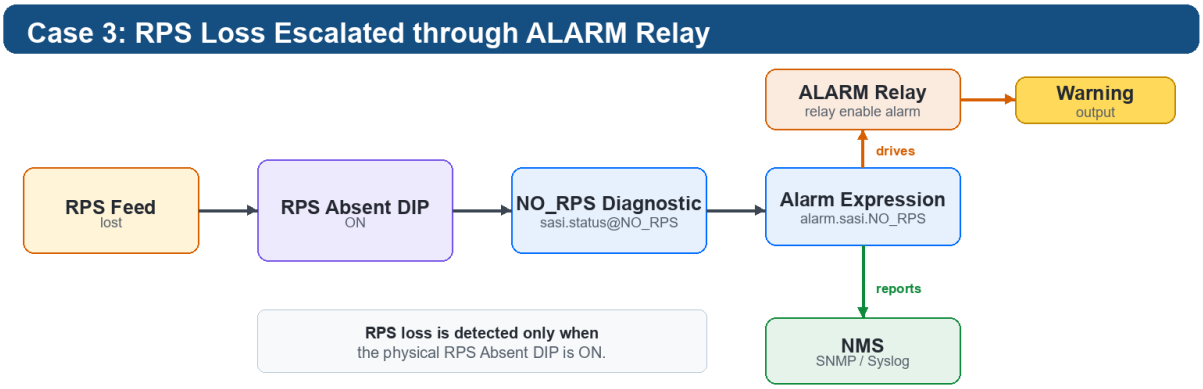

Figure: Case 3 - redundant-power (RPS) loss escalated through the ALARM relay.

4.5.6.4 Case 4 - Alarm cascade (uplink chain)

Multiple units can share an alarm chain: each unit reads the upstream relay on its DI1 and forwards the state through its DO to the next unit, with the last unit reporting to central monitoring. Keep local alarm rules on every unit so that a transient upstream blip does not cause downstream false alarms.

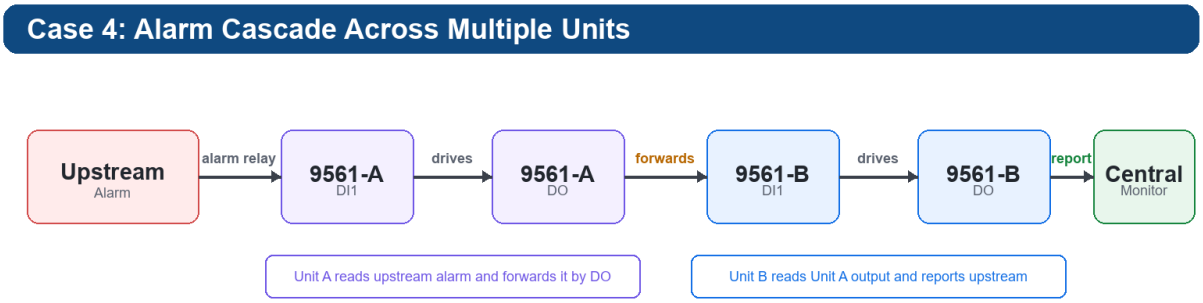

Figure: Case 4 - alarm cascade across multiple units.

4.5.7 Alarm Sources and Expressions

The alarm framework has three related concepts, summarised in the table below.

Concept	Quick description (with example)
Alarm source	An abnormal enumeration exposed by the switch. "show alarm sources" command can list all alarm source supported in system.
Monitored alarm source (named alarm / registered alarm)	The named alarm object created using the "alarm" command for a specific alarm source. System will periodically watch named alarm object. An abnormal event will set the named alarm object status.
Alarm expression	The Boolean expression defined on the named alarm object's status and will be evaluated periodically by system. When expression is evaluated as FALSE, it indicates an abnormal event raised. If expression is evaluated as TRUE, it indicates the named alarm object status is normal.

4.5.7.1 Alarm Source

An alarm source is an alarm enumeration the switch supports; show alarm sources lists them all. Different subsystems expose different abnormal events - a few examples by category:

sasi (digital I/O & environment)

Alarm source	Abnormal event
sasi.status@DI1	External contact on DI1 is abnormal
sasi.status@NO_PWR	Main power input lost
sasi.status@BOARD_T	Board temperature over threshold

Port

Alarm source	Abnormal event
port.status@Link	Port link down
port.status@LossOfSignal	SFP loss of signal
port.status@TxFault	SFP Tx fault

IP

Alarm source	Abnormal event

ip.status.acd.ipv4@ifIndex	IPv4 address conflict detected
ip.status.interface.ipv6@duplicated	Duplicate IPv6 address

MRP

Alarm source	Abnormal event
iecMrp.notification@RingOpen	MRP ring open
iecMrp.notification@MultipleMrms	Multiple MRP managers on the ring

4.5.7.2 Monitored Alarm Source (Named alarm)

A monitored alarm source (also called a named alarm or registered alarm) is created for a specific alarm source via the “alarm” command. Creating it tells the system to periodically watch that source's abnormal event.

alarm command syntax:

```
alarm <named_alarm> <expression>
```

Example - port link:

Alarm source	Named alarm
port.status@Link	alarm.port.link

```
configure terminal
 alarm alarm.port.link port.status["GigabitEthernet 1/1"]@Link==true
end
```

Example - DI1 / DI2:

Alarm source	Named alarm
sasi.status@DI1	alarm.sasi.DI1
sasi.status@DI2	alarm.sasi.DI2

```
configure terminal
 alarm alarm.sasi.DI1 sasi.status@DI1==false
 alarm alarm.sasi.DI2 sasi.status@DI2==false
end
```

If the alarm source is not supported, the command is rejected with "Monitoring object not defined".

4.5.7.3 Alarm Expression

An alarm expression is the boolean expression evaluated on a named alarm object status. Alarm expression is created in the “alarm” command which is used to create a named alarm object.

The alarm command syntax:

```
alarm <named_alarm> <expression>
```

Example: DI1 alarm command

```
configure terminal
 alarm alarm.sasi.DI1 sasi.status@DI1==false
end
```

This alarm command create a “**alarm.sasi.DI1**” named alarm object on “**sasi.status@DI1**” alarm source. The alarm expression is

```
sasi.status@DI1==false
```

After “**alarm.sasi.DI1**” named alarm object was created, system will periodically watch “**alarm.sasi.DI1**” status and evaluate the expression “**`sasi.status@DI1==false`**” based on named alarm’s status value. Named alarm status value is reading from the DI1 external contact value.

If DI1 external contact type is NO type, then it’s expression should be

```
sasi.status@DI1==false
```

If DI1 external contact type is NC type, then it’s expression should be

```
sasi.status@DI1==true
```

This is determined by following DI raw value table.

Contact type	Raw value	sasi.status@DI1	Alarm expression	Evaluate result	Alarm active
NO	0 (open)	false	sasi.status@DI1==false	true	false(Normal)
NO	1 (closed)	true	sasi.status@DI1==false	false	true(Abnormal)
NC	1 (closed)	true	sasi.status@DI1==true	true	false(Normal)
NC	0 (open)	false	sasi.status@DI1==true	false	true(Abnormal)

4.5.8 Alarm Escalation and Monitoring

The relay escalation path and the notification path are related but independent. The relay path is type-specific: DI alarms drive the DO relay path, while eligible system alarms such as PWR/RPS loss drive the ALARM relay path. The notification layer separately decides whether an active alarm is reported to syslog or as an SNMP trap. Each path is enabled separately.

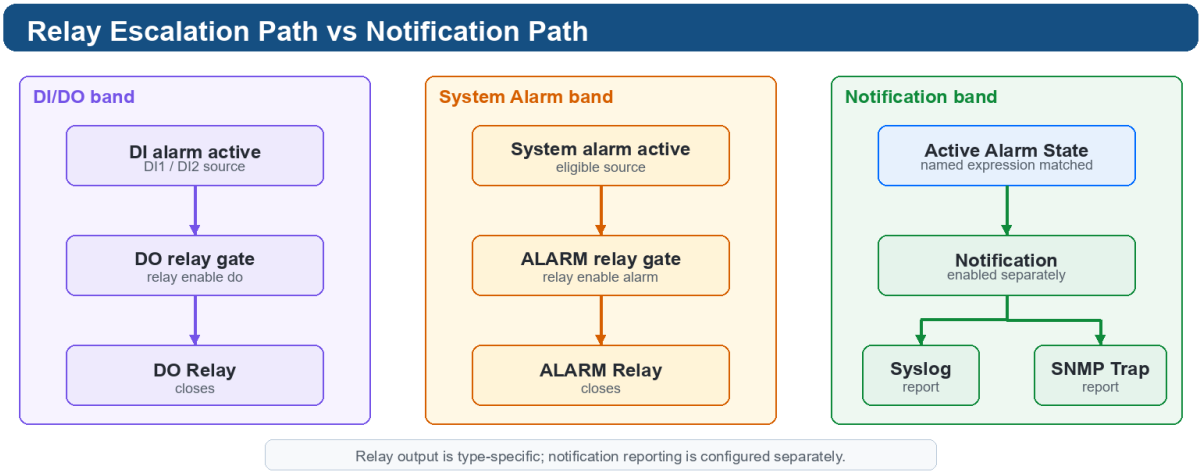

Figure: Relay escalation path and notification path are enabled separately.

Active Alarm: An Active alarm is a registered **named alarm** currently in the **Active state** - that is, its **expression evaluates FALSE**, so the monitored normal condition no longer holds. **Relay output** is type-specific: a DI active alarm uses the **DO relay** gate, while an eligible system active alarm uses the **ALARM relay** gate. **Syslog/SNMP notification** uses the active alarm state separately and still requires its own configuration. For the Alarm Expression concept, refer to Section 4.5.7.3 Alarm Expression; for ICLI command examples, refer to Section 4.5.8.3 ICLI Command Examples.

Before configuring syslog or SNMP trap notification, create a named Alarm Expression for the alarm source and verify the active or clear state with alarm status. The notification settings report the alarm state only; they do not control the physical DO or ALARM relay output.

4.5.8.1 Syslog Configuration

Syslog notification sends alarm active/clear events to the local log and, when a logging host is configured, to a remote syslog server. Subscribe to the alarm tree path alarm.status; do not use an individual alarm name as the notification node.

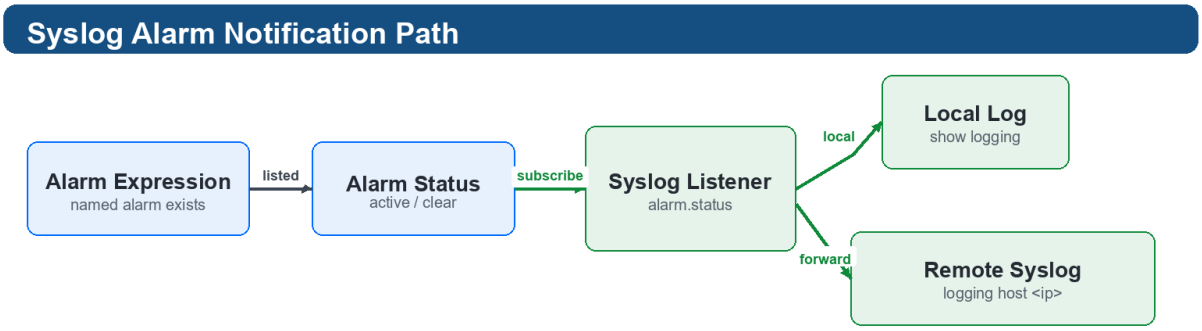

Figure: Syslog alarm notification path.

ICLI example:

```
configure terminal
 logging on
 logging level warning
 logging host <syslog-server-ip>
 logging notification listen alarmNotif level warning alarm.status
 end
show logging
```

The logging host forwards messages to the remote syslog collector. If no logging host is configured, the event is kept only in the local log.

4.5.8.2 SNMP Trap Configuration

SNMP trap notification sends alarmTrapStatus trap PDUs when a defined alarm changes state. Enable the alarm trap, configure a trap host, set the SNMP version and community, and keep the host entry enabled with no shutdown.

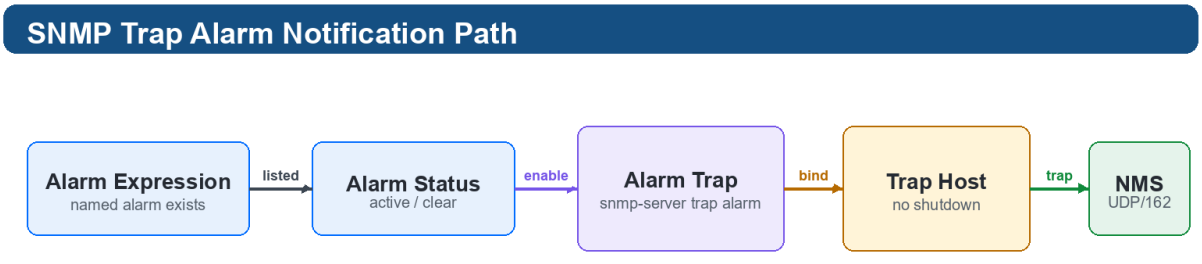

Figure: SNMP trap alarm notification path.

ICLI example:

```
configure terminal
 snmp-server trap alarmTrapStatus
```

```
 snmp-server host AlarmTrapHost
   host <nms-ip> 162 traps
   version v2 <community>
   no shutdown
   exit
 end
```

The alarm trap command may be expanded in running-config as alarmTrapStatus. Suppressed alarms are not reported to syslog or SNMP traps.

4.5.8.3 ICLI Command Examples

Basic status checks:

```
show dip status
show sasi alarms
show alarm sources
show alarm status
```

Configure the DI attach mode:

```
configure terminal
 di1 attachment normal-open
end
show di1 attachment
show di2 attachment
```

Create DI alarm expressions:

```
configure terminal
 alarm alarm.sasi.DI1 sasi.status@DI1==false
 alarm alarm.sasi.NO_RPS sasi.status@NO_RPS==false
end
```

Enable the relay outputs:

```
configure terminal
 relay enable do
 relay enable alarm
end
```

Create PWR / RPS alarm expressions:

```
configure terminal
 alarm alarm.sasi.NO_PWR sasi.status@NO_PWR==false
 alarm alarm.sasi.NO_RPS sasi.status@NO_RPS==false
end
```

Enable syslog alarm notification:

```
configure terminal
 logging on
 logging level warning
 logging host <server-ip>
 logging notification listen alarmNotif level warning alarm.status
end
```

Enable SNMP-trap alarm notification:

```
configure terminal
 snmp-server trap alarmTrapStatus
 snmp-server host AlarmTrapHost
  host <server-ip> 162 traps
  version v2 <community>
  no shutdown
  exit
```

```
end
```

4.5.9 DI event trigger and notification path

Before going live, simulate one event and confirm that the alarm is raised and then clears.

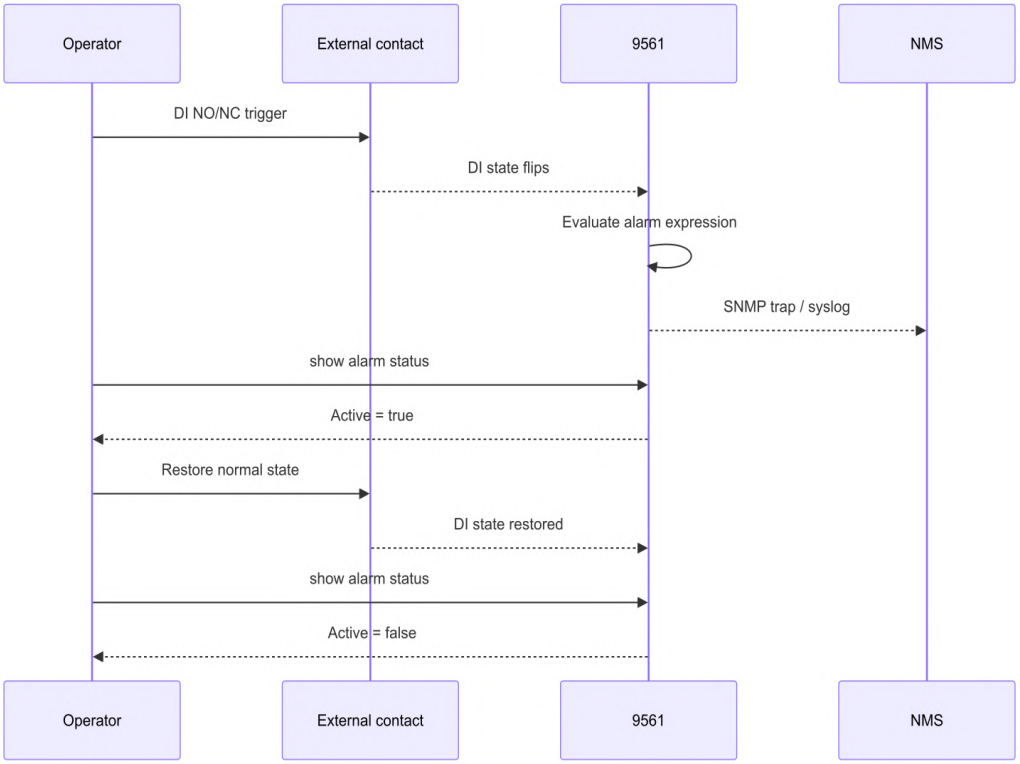

Figure: Recommended trigger-test sequence.

4.5.10 Alarm Suppression

Suppression mutes an alarm without deleting it. A suppressed alarm stays Active in show alarm status, but it no longer drives the DO or ALARM relay and sends no syslog or SNMP trap.
It affects only the output side, not the source value or the alarm state. Mute with alarm suppress <name>, restore with no alarm suppress <name>, delete with no alarm <name>. Rules survive a reboot only after copy running-config startup-config.
Use it to silence a known, accepted condition while keeping it visible in show alarm status.

Suppress an named alarm you have already handled (it stays Active but is muted):

```
configure terminal
 alarm suppress alarm.sasi.DI1
end
show alarm status
```

Restore the named alarm so its outputs resume:

```
configure terminal
 no alarm suppress alarm.sasi.DI1
end
```

Remove an alarm rule completely:

```
configure terminal
 no alarm alarm.sasi.DI1
end
```

Persist alarm rules across reboot:

```
copy running-config startup-config
```

4.6 CFM

CFM (Connectivity Fault Management) is a standard defined by IEEE802.1ag and is largely identical with ITU-T Recommendation Y.1731. It provides tools to detect, verify, and isolate Ethernet connectivity faults. It operates using maintenance domains, maintenance points, and standardized protocol messages to monitor continuity, identify fault locations, and support proactive network maintenance.

4.6.1 Global

CFM Global settings define system-wide parameters. These settings control overall CFM operation and must be configured before any maintenance domains or points become active.
CFM status can be viewed at Monitor→CFM.

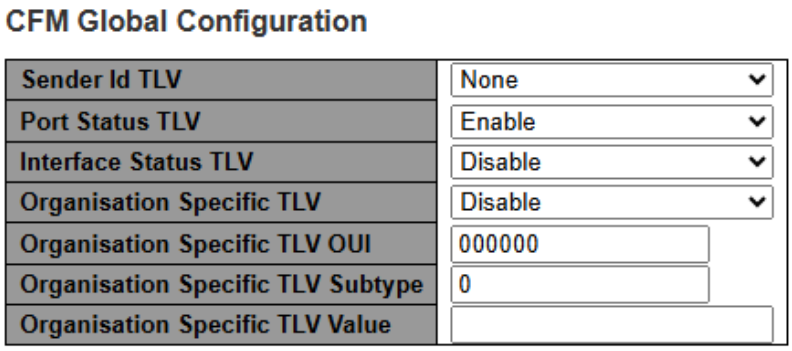

Parameter	Description
Sender Id TLV	Choose whether and what to use as Sender ID TLVs in CCMs generated by this switch. Can be overridden by Domain and Service level configuration. **None** **Chassis** **Manage** **ChassisManage**
Port Status TLV	Choose whether to send Port Status TLVs in CCMs generated by this switch. Can be overridden by Domain and Service level configuration. **Enable:** Send Port Status TLVs in CCMs generated by this switch. **Disable:** Do not send Port Status TLVs in CCMs generated by this switch.
Interface Status TLV	Choose whether to send Interface Status TLVs in CCMs generated by this switch. Can be overridden by Domain and Service level configuration. **Enable:** Send Interface Status TLVs in CCMs generated by this switch. **Disable:** Do not Send Interface Status TLVs in CCMs generated by this switch.
Organisation Specific TLV	Choose whether to send Organisation Specific TLVs in CCMs generated by this switch. Can be overridden by Domain and Service level configuration. **Enable:** Send Organisation Specific TLVs in CCMs generated by this switch. **Disable:** Do not send Organisation Specific TLVs in CCMs generated by this switch.
Organisation Specific TLV OUI	This is the three-bytes OUI transmitted with the Organization-Specific TLVs. Enter as 6 characters 0-9, a-f.
Organisation Specific TLV Subtype	This is the subtype transmitted with the Organization-Specific TLV. Can be any value in range **[0; 255]**
Organisation Specific TLV Value	This is the value transmitted in the Organization-Specific TLVs. Value is a printable character string of length 0-63.

4.6.2 Domain

CFM domain configuration defines maintenance domains and their levels, which represent administrative boundaries within the network. Each domain establishes the scope for fault management and determines which CFM messages are processed or forwarded.
CFM status can be viewed at Monitor→CFM.

CFM Domain Configuration

Delete	Domain	Format	Name	Level	TLV option select			
					Sender Id	Port Status	Interface Status	Org. Specific
*								
No entry exists								

Add New Entry

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
Domain	Name of Domain. Value is a single word which begins with an alphabetic letter A-Z or a-z with length 1-15.
Format	Select the MD name format. To mimic Y.1731 MEG IDs, use type None. **None** **String**
Name	The contents of this parameter depend on the value of the format member. If format is **None**: Name is not used, but will be set to all-zeros behind the scenes. This format is typically used by Y.1731-kind-of-PDUs. If format is **String**: Name must contain a string from 1 to 43 characters long.
Level	MD/MEG level of this domain. Valid values are restricted to **0 to 7**. **This switch exhibits Independent MEG level.** On Independent MEG level architectures, Port Down MEPs never perform level filtering on frames not classified to the MEP's VID. So if you have a Port MEP on VID X and a VLAN MEP on VID Y and an OAM frame ingresses any port on VID Z, it is not subject to handling/level-filtering by any of the two MEPs.
TLV option select	**Sender Id:** Default Sender ID TLV format to be used in CCMs generated by this Domain (may be overridden in service) **None:** Do not include Sender ID TLVs. **Chassis:** Enable Sender ID TLV and send Chassis ID (MAC Address). **Manage:** Enable Sender ID TLV and send Management address (IPv4 Address). **ChassisManage:** Enable Sender ID TLV and send both Chassis ID (MAC Address) and Management Address (IPv4 Address). **Defer:** Let the global configuration decide if Sender ID TLVs shall be included (may be overridden in service). **Port Status:** Include or exclude Port Status TLV in CCMs generated by this Domain or let higher level determine (may be overridden in Service). **Disable:** Do not include Port Status TLVs. **Enable:** Include Port Status TLVs. **Defer:** Let the global configuration decide if Port Status TLVs shall be included (may be overridden in Service). **Interface Status:** Include or exclude Interface Status TLV in CCMs generated by this Domain or let higher level determine (may be overridden in Service). **Disable:** Do not include Interface Status TLVs. **Enable:** Include Interface Status TLVs. **Defer:** Let the global configuration decide if Interface Status TLVs shall be included (may be overridden in Service). **Org. Specific:** Exclude Organization-Specific TLV in CCMs generated by this Domain or let higher level determine (may be overridden in Service). **Disable:** Do not include Organization-Specific TLVs.

	Defer: Let the global configuration decide if Organization-Specific TLVs shall be included (may be overridden in Service).

4.6.3 Service

CFM Service configuration defines maintenance associations and endpoints for specific Ethernet services. It specifies which interfaces participate, the roles of maintenance points (MIPs or MEPs), and the monitoring parameters for continuity and fault detection.
CFM status can be viewed at Monitor→CFM.

CFM Service Configuration

Delete	Domain	Service	Format	Name	VLAN	CCM Interval	TLV option select			
							Sender Id	Port Status	Interface Status	Org. Specific
*										
No entry exists										

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
Domain	Name of Domain under which this Service resides.
Service	Name of Service. Value is a single word which begins with an alphabetic letter A-Z or a-z with length 1-15.
Format	Select the short Service name format. This decides how the value of the Name parameter will be interpreted. To mimic Y.1731 MEG IDs, create an MD instance with an empty name and use Y1731 ICC or Y1731 ICC CC. Possible values are: **String**, **Two Octets**, **Y1731 ICC**, **Y1731 ICC CC** Look under *Name* for explanation.
Name	The contents of this parameter depend on the value of the format member. Besides the limitations explained for each of them, the following applies in general: If the Domain *Format* is **None**, the size of this cannot exceed 45 bytes. If the Domain *Format* is not **None**, the size of this cannot exceed 44 bytes. If *Format* is **String**, the following applies: length must be in range [1; 44] Contents must be in range [32; 126] If *Format* is **Two Octets**, the following applies: *Name[0]* and *Name[1]* will both be interpreted as unsigned 8-bit integers (allowing a range of [0; 255]). *Name[0]* will be placed in the PDU before *Name[1]*. The remaining available bytes in name will not be used. If *Format* is **Y1731 ICC**, the following applies: length must be 13. Contents must be in range [a-z,A-Z,0-9] Y.1731 specifies that it is a concatenation of ICC (ITU Carrier Code) and UMC (Unique MEG ID Code): ICC: 1-6 bytes UMC: 7-12 bytes In principle UMC can be any value in range [1; 127], but this API does not allow for specifying length of ICC, so the underlying code doesn't know where ICC ends and UMC starts. The Domain *Format* must be **None**. If *Format* is **Y1731 ICC CC**, the following applies: length must be 15. First 2 chars (CC): Must be amongst [A-Z] Next 1-6 chars (ICC): Must be amongst [a-z,A-Z,0-9]

	Next 7-12 chars (UMC): Must be amongst [a-z,A-Z,0-9] There may be ONE (slash) present in name[3-7]. The Domain *Format* must be **None**.
VLAN	The MA's primary VID. A primary VID of 0 means that all MEPs created within this MA will be created as port MEPs (interface MEPs). There can only be one port MEP per interface. A given port MEP may still be created with tags, if that MEP's VLAN is non-zero." A non-zero primary VID means that all MEPs created within this MA will be created as VLAN MEPs. A given MEP may be configured with another VLAN than the MA's primary VID, but it is impossible to have untagged VLAN MEPs.
CCM Interval	The CCM rate of all MEPs bound to this Service.
TLV option select	**Sender Id:** Default Sender ID TLV format to be used in CCMs generated by this Service. **None:** Do not include Sender ID TLVs. **Chassis:** Enable Sender ID TLV and send Chassis ID (MAC Address). **Manage:** Enable Sender ID TLV and send Management address (IPv4 Address). **ChassisManage:** Enable Sender ID TLV and send both Chassis ID (MAC Address) and Management Address (IPv4 Address). **Defer:** Let the Domain configuration decide if Sender ID TLVs shall be included. **Port Status:** Include or exclude Port Status TLV in CCMs generated by this Service or let higher level determine. **Disable:** Do not include Port Status TLVs. **Enable:** Include Port Status TLVs. Defer Let the Domain configuration decide if Port Status TLVs shall be included. **Interface Status:** Include or exclude Interface Status TLV in CCMs generated by this Service or let higher level determine. **Disable:** Do not include Interface Status TLVs. **Enable:** Include Interface Status TLVs. **Defer:** Let the Domain configuration decide if Interface Status TLVs shall be included. **Org. Specific:** Exclude Organization-Specific TLV in CCMs generated by this Service or let higher level determine. **Disable:** Do not include Organization-Specific TLVs. **Defer:** Let the Domain configuration decide if Organization-Specific TLVs shall be included.

4.6.4 MEP

A Maintenance End Point (MEP) is a CFM maintenance point located at the boundary of a Maintenance Domain. MEPs generate and respond to CFM protocol messages, such as Continuity Check Messages (CCM), Loopback Messages (LBM), and Link trace Messages (LTM), to monitor service integrity and detect connectivity faults.
This switch supports two types of down MEPs: port-based down MEPs and VLAN-based down MEPs.
CFM MEP Status is displayed at Monitor→CFM→Status.

Port Down-MEPs

In 802.1Q terminology, Port MEPs are located below the EISS entity, that is, closest to the physical port. Port MEPs are used by e.g. APS for protection purposes.
Port MEPs are created when the encompassing service has type "Port".
Port MEPs may send OAM PDUs tagged or untagged. An OAM PDU will be sent untagged only if the MEP's VLAN is set to "Inherit" (0). Any other value will cause it to be sent tagged with the port's TPID, whether or not the VLAN matches the port's PVID and that PVID is meant to be sent untagged.

VLAN Down-MEPs

in 802.1Q terminology, VLAN MEPs are located above the EISS entity.
This means that tagging of OAM PDUs will follow the port's VLAN configuration.
Thus, if a VLAN MEP is created on the Port's PVID and PVID is configured to be untagged, OAM PDUs will be transmitted untagged.

VLAN MEPs are created when the encompassing service has type "VLAN".

Down-MEP creation rules

There are a few rules to obey when creating Down-MEPs:

1. There can only be one Port MEP on the same port.
2. There can only be one VLAN MEP on the same port and VLAN.
3. A VLAN MEP must have a higher MD/MEG level than a Port MEP on the same port and VLAN.

These checks are performed automatically on administratively enabled MEPs when you change a particular MEP, change the Service Type from Port to VLAN or vice versa, or change the domain's MD/MEG level.

Port Down-MEPs

In 802.1Q terminology, Port MEPs are located below the EISS entity, that is, closest to the physical port. Port MEPs are used by e.g. APS for protection purposes.

Port MEPs are created when the encompassing service has type "Port".

Port MEPs may send OAM PDUs tagged or untagged. An OAM PDU will be sent untagged only if the MEP's VLAN is set to "Inherit" (0). Any other value will cause it to be sent tagged with the port's TPID, whether or not the VLAN matches the port's PVID and that PVID is meant to be sent untagged.

VLAN Down-MEPs

in 802.1Q terminology, VLAN MEPs are located above the EISS entity.

This means that tagging of OAM PDUs will follow the port's VLAN configuration.

Thus, if a VLAN MEP is created on the Port's PVID and PVID is configured to be untagged, OAM PDUs will be transmitted untagged.

VLAN MEPs are created when the encompassing service has type "VLAN".

Down-MEP creation rules

There are a few rules to obey when creating Down-MEPs:

1. There can only be one Port MEP on the same port.
2. There can only be one VLAN MEP on the same port and VLAN.
3. A VLAN MEP must have a higher MD/MEG level than a Port MEP on the same port and VLAN.

These checks are performed automatically on administratively enabled MEPs when you change a particular MEP, change the Service Type from Port to VLAN or vice versa, or change the domain's MD/MEG level.

CFM Mep Configuration

Delete	Domain	Service	MEPID	Direction	Port	VLAN	PCP	SMAC	Alarm Control			State Control		Remote MEPID
									Level	Present	Absent	CCM	Admin	
*														
No entry exists														

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
Domain	Name of Domain under which this MEP resides.
Service	Name of Service under which this MEP resides.
MEPID	The identification of this MEP. Must be an integer **[1..8091]**
Direction	Set whether this MEP is an Up- or a Down-MEP.
Port	Port on which this MEP resides.
VLAN	VLAN ID. Use the value 0 to indicate untagged traffic (implies a port MEP)..

PCP	Choose PCP value in PDUs' VLAN tag. Not used if untagged.
SMAC	Set a Source MAC address to be used in CCM PDUs originating at this MEP. Must be a unicast address. Format is **XX:XX:XX:XX:XX:XX**. If all-zeros, the switch port's MAC address will be used instead.
Alarm Control	**Level:** If a defect is detected with a priority higher than this level, a fault alarm notification will be generated. Valid range is **[1; 6]** with 1 indicating that any defect will cause a fault alarm and 6 indicating that no defect can cause a fault alarm. See 802.1Q-2018, clause 20.9.5, LowestAlarmPri The possible defects and their priorities are: **Short name** / **Description** / **Priority** **DefRDICCM** Remote Defect Indication 1 **DefMACstatus** MAC Status 2 **DefRemoteCCM** Remote CCM 3 **DefErrorCCM** Error CCM Received 4 **DefXconCCM** Cross Connect CCM Received 5 **Present:** The time in milliseconds that defects must be present before a fault alarm notification is issued. Default is 2500 ms. **Absent:** The time in milliseconds that defects must be absent before a fault alarm notification is reset. Default is 10000 ms.
State Control	**CCM:** Enable or disable generation of continuity-check messages (CCMs) **Admin:** Enable or disable this MEP. When this MEP is enabled, it will check received/missing CCMs and can raise defects.
Remote MEPID	Specify the Remote MEP that this MEP is expected to receive CCM PDUs from. Must be an integer **[0..8091]** where 0 means undefined. The value of Remote MEPID must be different from the value of MEPID.

4.7 APS

The Automatic Protection Switching (APS) function provides fast failover for Ethernet services by automatically switching traffic from a failed working path to a preconfigured protection path. It implements a ring protection protocol and linear protection mechanisms for point-to-point, VLAN-based Ethernet Sub-Network Connections (ETH SNC) in carrier and industrial Ethernet transport networks. APS is defined by the ITU-T G.8031 standard and enables rapid service recovery in case of link or node failures, helping to minimize packet loss and service disruption.

The APS status can be checked under Monitor→APS.

APS Configuration

APS #	Mode	SMAC	Level	VLAN	PCP	Rev	TxAps	WTR	HoldOff	Enable
0	1:1	00:00:00:00:00:00	0	0	7	☐	☐	300	0	☐

Parameter	Description
APS #	The ID of the APS. Maximum number of creatable APS instances is 12. The ID links to the APS instance page, where you can reset counters and issue commands.
Mode	**1:1** This will create a 1:1 APS. In the linear 1:1 protection switching architecture, the protection transport entity is dedicated to the working transport entity. However, the normal traffic is transported either on the working transport entity or on the protection transport entity using a selector bridge at the source of the protected domain. The selector at the sink of the protected domain selects the entity which carries the normal traffic. **1+1 Uni** This will create a 1+1 Unidirectional APS. **1+1 Bi** This will create a 1+1 Bidirectional APS.

	In the linear 1+1 protection switching architecture, a protection transport entity is dedicated to each working transport entity. The normal traffic is copied and fed to both working and protection transport entities with a permanent bridge at the source of the protected domain. The traffic on working and protection transport entities is transmitted simultaneously to the sink of the protected domain, where a selection between the working and protection transport entities is made based on some predetermined criteria, such as server defect indication.
Level	MD/MEG Level (0-7).
SMAC	Set a Source MAC address to be used in L-APS PDUs. Must be a unicast address. Format is **XX:XX:XX:XX:XX:XX**. If all-zeros, the switch port's MAC address will be used instead.
VLAN	The VLAN ID used in the L-APS PDUs. 0 means untagged.
PCP	PCP (priority) (default 7). The PCP value used in the VLAN tag unless the L-APS PDU is untagged. Must be a value in range **0 - 7**.
Rev	When checked, the port recovery mode is revertive, that is, traffic switches back to the working port after the condition(s) causing a switch has cleared. In the case of clearing a command (e.g. forced switch), this happens immediately. In the case of clearing of a defect, this generally happens after the expiry of the WTR (Wait-To-Restore) timer. When unchecked, the port recovery mode is non-revertive and traffic is allowed to remain on the protect port after a switch reason has cleared.
TxAps	Choose whether this end transmits APS PDUs. Only used for 1+1, unidirectional.
WTR	When Rev is checked, WTR (Wait-To-Restore) tells how many seconds to wait before restoring to the working port after a fault condition has cleared. Valid range **1 - 720**
HoldOff	When a new (or more severe) defect occurs, the hold-off timer will be started and the event will be reported after the timer expires. HoldOff time is measured in milliseconds, and valid values are in the range 0 - 10000. Default is 0, which means immediate reporting of the defect.
Enable	The administrative state of this APS instance. Check to make it function normally and uncheck to make it cease functioning.

The paths for *Working* and *Protecting* states have to be assigned to a port on the switch. If **Down-MEP** is used as *SF Type*, the *SF MEP* instance (**Domain::Service::MEPID**) representing the working flow has to be defined. The selected MEP instance does not need to exist when this APS is configured.

APS Signal Fail Trigger

Working					Protecting				
Port	SF Type	Domain	Service	MEPID	Port	SF Type	Domain	Service	MEPID
1	Link			0	1	Link			0

Parameter	Description
Port	The Port this flow is attached to.
SF Type	Selects whether Signal Fail (SF) comes from the link state of a given Port, or from a Down-MEP.
Domain	The domain of the SF MEP.
Service	The service of the SF MEP.
MEPID	The MEPID of the SF MEP.

4.8 ERPS

The ITU-T G.8032 **Ethernet Ring Protection Switching (ERPS)** feature implements protection switching mechanisms to provide protection for Ethernet traffic in a ring topology, while ensuring that no loops are

within the ring at the Ethernet layer. The loops are prevented by blocking traffic on either a predetermined link or a failed link.

The Ethernet ring protection functionality includes the following:

- Loop avoidance
- The use of learning, forwarding, and Filtering Database (FDB) mechanisms

Loop avoidance in an Ethernet ring is achieved by guaranteeing that, at any time, traffic may flow on all but one of the ring links. This particular link is called the **ring protection link (RPL)** and under normal conditions this ring link is blocked, i.e., not used for service traffic. One designated Ethernet ring node, the **RPL owner** node, is responsible to block traffic at one end of the RPL. Under an Ethernet ring failure condition, the RPL owner node is responsible for unblocking its end of the RPL, unless the RPL has failed, allowing the RPL to be used for traffic. The other Ethernet ring node adjacent to the RPL, the **RPL neighbor** node, may also participate in blocking or unblocking its end of the RPL.

The Ethernet rings could support a multi-ring/ladder network that consists of conjoined Ethernet rings by one or more interconnection points. The protection switching mechanisms and protocol defined in this recommendation shall be applicable for a multi-ring/ladder network, if the following principles are adhered to:

- R-APS channels are not shared across Ethernet ring interconnections;
- on each ring port, each traffic channel and each R-APS channel are controlled (e.g., for blocking or flushing) by the Ethernet ring protection control process (ERP control process) of only one Ethernet ring;
- Each major ring or sub-ring must have its own RPL.

In an Ethernet ring, without congestion, with all Ethernet ring nodes in the idle state (i.e., no detected failure, no active automatic or external command and receiving only "NR, RB" R-APS messages), with less than 1200 km of ring fiber circumference and fewer than 16 Ethernet ring nodes, the switch completion time (transfer time as defined in [ITU-T G.808.1]) for a failure on a ring link shall be less than **50ms**.

The status of all ERPS rings can be viewed at Monitor→ERPS.

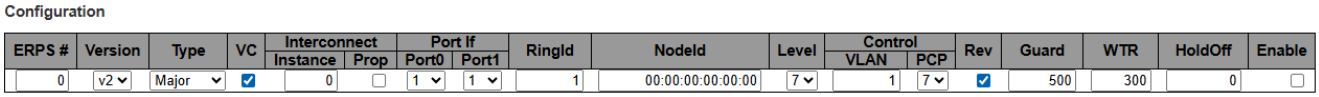
Configuration

ERPS #	Version	Type	VC	Interconnect		Port If		RingId	NodeId	Level	Control		Rev	Guard	WTR	HoldOff	Enable
				Instance	Prop	Port0	Port1				VLAN	PCP					
0	v2	Major	☑	0	☐	1	1	1	00:00:00:00:00:00	7	1	7	☑	500	300	0	☐

Parameter	Description
ERPS #	The ID of ERPS. The allowed value is from **1 - 64**.
Version	ERPS protocol version. v1 and v2 are supported.
Type	Type of ring. Possible values: **Major**: ERPS major ring (G.8001-2016, clause 3.2.39) **Sub**: ERPS sub-ring (G.8001-2016, clause 3.2.66) **InterSub**: ERPS sub-ring on an interconnection node (G.8001-2016, clause 3.2.66)
VC	Controls whether to use a Virtual Channel with a sub-ring.
Interconnect Instance	For a sub-ring on an interconnection node, this must reference the instance ID of the ring to which this sub-ring is connected.
Interconnect Prop	Controls whether the ring referenced by Interconnect Instance shall propagate R-APS flush PDUs whenever this sub-ring's topology changes.
Ring Id	The Ring ID is used - along with the control VLAN - to identify R-APS PDUs as belonging to a particular ring.

Node Id	The Node ID is used inside the R-APS specific PDU to uniquely identify this node (switch) on the ring.
Level	MD/MEG Level of R-APS PDUs we transmit.
Control PCP	The PCP value used in the VLAN tag of the R-APS PDUs.
Rev	If ticked, revertive mode restores the failed link when it recovers and blocks the RPL. Non-revertive mode keeps using the RPL and blocks the failed link instead when it recovers.
Guard	After a switchover event, the Guard timer blocks incoming requests for the specified time to avoid oscillation and ensure ring stability. Valid range is **10 - 2000** ms.
WTR	The Wait-to-Restore timer delays switching traffic back to the recovered working path after a fault clears. It ensures the path is stable before restoration, preventing unnecessary or repeated switching. Valid range 1 - 720 sec.
Hold Off	The Hold-off Timer delays fault detection or protection switching for a short, configurable period before switching. It filters transient or short-duration faults to prevent unnecessary protection actions. Value is rounded down to 100ms precision. Valid range is **0 - 10000** ms.
Enable	The administrative state of this ERPS. Check to make it function normally and uncheck to make it cease functioning.

Signal Fail Trigger

Port0				Port1			
Type	Domain	Service	MEPID	Type	Domain	Service	MEPID
Link			0	Link			0

Parameter	Description
Type	Selects whether Signal Fail (SF) comes from the link state of a given interface, or from a Down-MEP.
Domain, Service, MEPID	Identification of the MEP instance to provide Signal Fail, if Type is MEP.

Protected VLANs

VLAN ID	

Ring Protection Link

RPL Mode	RPL Port
None	RingPort0

Parameter	Description
VLAN ID	VLANs which are protected by this ring instance. At least one VLAN must be protected. Specify as a comma separated list of VLAN IDs or VLAN ranges. Ex.: 1,4,7,30-70
RPL Mode	Ring Protection Link mode. One of **None**: This switch doesn't have the RPL port in the ring **Owner**: This switch is RPL owner for the ring (G.8001-2016, clause 3.2.61) **Neighbor**: This switch is RPL neighbor for the ring (G.8001-2016, clause 3.2.60)
RPL Port	Indicates whether it is *Port0* or *Port1* that is the Ring Protection Link. Not used if RPL Mode is **None**.

4.9 Media Redundancy Protocol

The **Media Redundancy Protocol (MRP)** defined in IEC 62439-2 is a recovery protocol based on a ring topology, designed to react deterministically on a single failure of an inter-switch link or switch in the network and restore traffic within a predictable time window under the control of a dedicated media redundancy node.
Every ring requires one Media Redundancy Manager (MRM) that monitors the control VLAN, evaluates signal-fail triggers, and blocks or unblocks the two ring ports; the remaining nodes operate as Media Redundancy

Clients (MRCs) and may include an Auto Manager (MRA) that steps in if the manager vanishes. To join two separate rings, selected nodes are elevated to Media Redundancy Interconnection Managers (MIMs) or Clients (MICs) so a third interconnection ring can carry alternate paths.

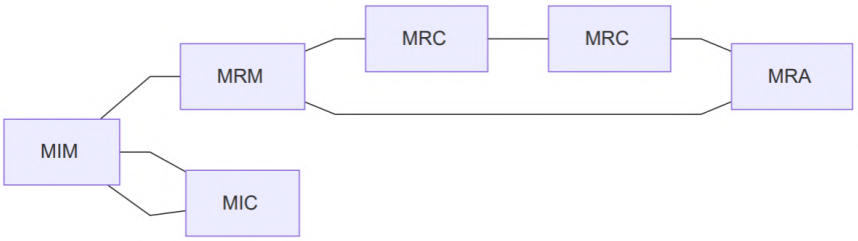

Media Redundancy Protocol status can be viewed at Monitor→Media Redundancy.

MRP Configuration Overview

This *MRP Configuration Overview* shows the defined MRP instances and allows for creating new and editing or deleting existing.
The meaning of the fields is described in detail in the *MRP Detailed Configuration* section below.

Media Redundancy Protocol Configuration Overview

Auto-refresh ☐ Refresh

Inst #	Enabled	Ring									Interconnection						Operational		
		Role	Name	Port1 Port	Port1 SF	Port2 Port	Port2 SF	VLAN	Recovery Profile	UUID	Role	Name	Port	SF	VLAN	Recovery Profile	State	Warnings	
1	✕	Auto Manager		none	Link	none	Link	0	500 ms	ffffffffffffffffffffffffffffffff	None		none	Link	0	500 ms			

Parameter	Description
Inst #	The MRP instance ID. Click on the instance ID to see status and statistics for it.
Enabled	Checked if this MRP instance is administratively enabled. Unchecked otherwise. *Note: Even if enabled, it may not be functioning properly. See Operational State and Operational Warnings columns for this.*
Ring Role	Shows the configured MRP ring role.
Ring Name	Shows the configured name of this MRP instance.
Ring Port1/2 Port	Shows the configured ring port interfaces of this MRP instance. none if not configured.
Ring Port1/2 SF	Shows the configured signal fail trigger mechanism for the ring port interfaces.
RING VLAN	Shows the VLAN ID that MRP_Xxx PDUs are transmitted with. 0 if transmitting untagged.
RING Recovery Profile	Shows the configured recovery profile for this MRP instance.
RING UUID	Shows the configured UUID/DomainID for this MRP instance.
Interconnection Role	Shows the configured interconnection role this MRP instance.
Interconnection Name	Shows the configured interconnection name of this MRP instance.
Interconnection Port	Shows the configured interconnection port interface of this MRP instance. 'none' if not configured.
Interconnection SF	Shows the configured signal fail trigger mechanism for the interconnection port interface.

Interconnection VLAN	Shows the VLAN ID that MRP_InXxx PDUs are transmitted with. 0 if transmitting untagged.
Interconnection Recovery Profile	Shows the configured interconnection recovery profile for this MRP instance.
Enable	Checkbox to administratively enable or disable the MRP instance. Disabling an instance allows you to edit its parameters or temporarily unblock the ring.
Operational State	Reflects the operational state for this instance: `Green`: Active. `Red`: Disabled or recovering *Note: Even if the instance is operational, it may have operational warnings. See next column description.*
Operational Warnings	Indicates if operational warnings (OperWarnings) are present: `Yellow`: Warnings exist (e.g., VLAN mismatch, signal fail fallback, multiple managers). Hover the mouse over the yellow indicator to view detailed warning text. `Gray`: No warnings or the instance is not enabled.

MRP Rings

Ring Ports

Each ring instance has two dedicated ports known as ring ports. These are often referred to as Port1 and Port2. Any port on the switch can be assigned as a ring port, but the standard mandates the link speed to be at least 100 Mbps, full duplex.

Note: This is not checked by the MRP implementation.

MRP Roles

MRM

An MRP ring must have one Media Redundancy Manager (MRM), which is responsible for controlling the ring state. Initially, the MRM blocks one ring port and unblocks the other. When a ring port is blocked, no frames can ingress and egress that port except for certain CPU-destined and -generated frames. The unblocked (forwarding) ring port is called the primary ring port. The other is called the secondary ring port.
The MRM sends MRP_Test PDUs at a configured time period in both directions of the ring and if it does not receive its own PDUs after a certain time, it unblocks the secondary ring port and transmits an MRP_TopologyChange PDU, which contains an MRP_Interval parameter that indicates when the remaining nodes of the ring (the so-called Media Redundancy Clients (MRCs) must flush their forwarding database (FDB).

Note: Only one node in the ring may be configured as an MRM. If you have multiple nodes configured as MRM, one node will block for the other node's MRP_Test PDUs. As a consequence, each MRM will unblock both its ring ports, but still not pass MRP_Test PDUs through from one ring port to the other. This is likely to cause loops in the network.

MRC

MRCs will - in normal operation - unblock both its ring ports and allow MRP_Test PDUs from the MRM to pass through. Upon link failure, these PDUs cannot pass through and the MRM detects the failure and unblocks its blocked ring port.
On link change, the MRC transmits an MRP_LinkChange PDU, which can take the form of an MRP_LinkUp or MRP_LinkDown PDU.
An MRC is responsible for forwarding all MRP PDUs from one ring port to the other - if possible.

Note: If all nodes in the ring are configured as MRCs, there is no MRM to block one of its ring ports, and since an MRC by default unblocks both of its ring ports, it is likely to cause loops in the network.

MRA

The last MRP role is the Media Redundancy Automanager (MRA), which can take the form of either an MRC or an MRM.
Instead of configuring a node as an MRM or an MRC, it may be configured as an MRA. MRAs select one MRM among each other by using a voting protocol based on MRP_Test PDUs and MRP_TestMgrNAck PDUs.
Each MRP_Test PDU contains a priority field, which is provisioned by the end user. A lower numerical value indicates a higher priority.
Whenever a new MRA joins the ring, it starts sending MRP_Test PDUs with its own priority. All MRAs receive these MRP_test PDUs and compare the priority field to their own priority.
If an MRA's own priority is higher (numerically lower), it returns an MRA_MgrTestNAck PDU to the sender and remains an MRM. The recepient of an MRA_Test_MgrNAck PDU ceases its MRM role and becomes an MRC.
If an MRA's own priority is lower (numerically higher), it becomes an MRC.
If an MRA's own priority is equal to the received priority, it's the node's MAC address that determines the behavior: The lower MAC address, the higher priority.
MRAs operating as MRCs continuously monitor MRP_Test PDUs to see if they need to return to the MRM role in case of missing MRP_Test PDUs or in case it has higher priority than what it receives.

Note: It must be ensured that a given ring does not contain nodes configured as MRMs as well as nodes configured as MRAs. The reason is that a node configured as an MRM doesn't have the ability to send MRP_TestMgrNAck PDUs back to a node configured as MRA (even though the standard specifies that the priority of an MRM must be numerically lower than the priority of an MRA). An MRA will therefore think that it is the manager of the ring, causing the ring to have multiple MRMs, which may cause loops. So in short, the priority of a node is only relevant for MRAs and not MRMs.

Recovery Profiles
The MRM and MRC contain a set of parameters that specify the maximum recovery time of a ring. Table 59 and Table 60 of the standard outline four consistent sets of such maximum recovery times for MRM and MRC, respectively. These are 10 ms, 30 ms, 200 ms, and 500 ms. Here, one set of parameters is called a recovery profile.

Note: Not all four recovery profiles are available on all chips. Only the two slowest (200 and 500 ms) are available on software-based MRP implementations, whereas also the two fastest (10 ms and 30 ms profiles) are available on hardware-based MRP implementations.

Note: It is very important that all nodes in a ring run the same recovery profile. If not, MRAs may repeatedly turn from MRC to MRM and back to MRC, making the ring unstable.

Signal Fail Trigger
The MRP implementation supports either the physical link state on the ring ports to detect signal fail and recovery or the use of CFM MEPs.
If the ring ports are connected back-to-back with its partner node (which is a normal case, because otherwise MRP_Xxx PDUs will flood the network unless otherwise handled by the intermediate switches, e.g. through VLAN configuration), you may rely on using link state.

Note: The recommended CCM rates of 3.3 ms or 10 ms might not be supported by all switches. In that case, it makes sense to only use the link state and not CFM MEPs as signal fail trigger.

Interconnected Rings

The standard provides a method for interconnecting rings. To redundantly connect two MRP rings, two nodes of each ring are assigned additional roles, besides being operating as either MRM or MRC.

One of the nodes is assigned the role of a Media redundancy Interconnection Manager (MIM). The remaining three nodes (one in the MIM's ring and the two of the other ring) are assigned the role of Media redundancy Interconnection Clients (MICs).
The MIM connects to a MIC in the other ring, and the two remaining MICs connect to each other. In this way, a third ring - the interconnection ring - is formed between the four interconnection (I/C) nodes.
It is possible to interconnect more than two rings in a ladder topology. To interconnect N rings, N-1 interconnection rings are required.

Each interconnection ring must have its own, unique ID, called the interconnection ID, or IID in short.

MIM
The function of the MIM is to observe and control the redundant I/C topology. It does so by blocking and unblocking the I/C port as it sees fit. The MIM operates in one of two modes:

`Link Check`
In link check (LC) mode, the MIM starts by sending MRP_InLinkStatusPoll PDUs to ask for the MICs' current link status. The MICs in turn reply with MRP_InLinkUp if their I/C port has link or MRP_InLinkDown PDUs if their I/C port has signal fail.
A MIC autonomously sends MRP_InLinkUp or MRP_InLinkDown PDUs whenever the I/C port changes its link status, and the MIM blocks or unblocks its I/C port.

`Ring Check`
In ring check (RC) mode, the MIM continuously transmits MRP_InTest PDUs on its I/C port. In addition, if it's an MRC, it sends to both it ring ports, and if it's an MRM it sends to forwarding ring ports. If these PDUs come back to the MIM, the interconnection ring is considered closed, and the MIM blocks its I/C port. If they don't, the MIM considers the interconnection ring open, and the MIM unblocks its I/C port.

No matter the interconnection mode, the MIM transmits MRP_InTopologyChange PDUs whenever it changes its I/C port's forwarding state. These PDUs ask the three MICs to flush their FDB at a certain point in time.
An MRM picks up these PDUs and transforms them into MRP_TopologyChange PDU and transmits them on both its ring ports, asking the MRCs to also flush their FDB.

Note: The switches support both interconnection modes, but if more than two rings are to be interconnected, the Link Check mode must be used in all interconnection rings.

MIC
The MIC is responsible for forwarding MRP_InXxx PDUs between its ring and I/C ports and react on MRP_InTopologyChange and MRP_InLinkStatusPoll PDUs from the interconnection's MIM.

The MIC is also configured with either a Link Check or Ring Check mode. The main difference is that the MIC only replies to MRP_InLinkStatusPoll PDUs from the MIM in Link Check mode and that it only forwards MRP_InTest PDUs between ring and I/C ports in Ring Check mode.

Recovery Profiles
The MIM and MIC contain a set of parameters that specify the maximum recovery time of an interconnection. Table 61 and Table 62 of the standard outline two consistent sets of such maximum recovery times for MIM and MIC, respectively. These are 200 ms, and 500 ms. Here, one set of parameters is called a recovery profile.

Note: It is important that all MIM and MIC nodes in the same interconnection run the same recovery profile.

Signal Fail Trigger

The standard mandates CFM MEP instances be instantiated in both ends of the interconnection link when the interconnection topology is in LC-mode. It may use MEPs when the interconnection is in RC-mode. Note: This implementation allows for using link state as signal fail trigger in both LC and RC mode. It is up to the end user to determine a suitable mode.

Note: The standard suggests a very specific coding of a MEP's Maintenance Association ID (MAID). This coding suggests a dynamic change of the Maintenance Domain Name depending on the state of MRP.
Since CFM and MRP are decoupled entities in this implementation, it is not possible to have MRP dynamically change the Maintenance Domain Name.
The end user may configure the MEPs in any way the CFM module supports, one of which may be as close as possible to what the MRP standard suggests.

MRP Configuration

The *MRP Configuration* page appears when clicking the ⓔ or ⊕ symbol and allows for editing existing or new MRP instances. It is divided into Ring Configuration and Interconnection Configuration, each with its own Signal Fail (SF) trigger settings.

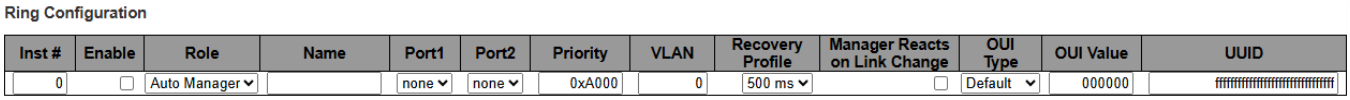
Ring Configuration

Inst #	Enable	Role	Name	Port1	Port2	Priority	VLAN	Recovery Profile	Manager Reacts on Link Change	OUI Type	OUI Value	UUID
0	☐	Auto Manager		none	none	0xA000	0	500 ms	☐	Default	000000	ffffffffffffffffffffffffffffffff

Parameter	Description
Inst #	The ID of the MRP instance. The allowed value is from **1 - 15**. If editing an existing instance, this cannot be changed.
Enable	Once all configuration parameters (including the Interconnection Configuration) is in place, check this to enable the MRP instance. Uncheck to disable the MRP instance and remove all configuration from hardware while clearing status and statistics. *Note: Be aware that disabling an MRP instance may cause loops in the network if not properly guarded by other means (e.g. Spanning Tree Protocol (STP)).*
Role	Configures the node's role in the ring: `Client (MRC)`: The node acts as a Media Redundancy Client. It forwards frames and blocks ring ports as instructed by the Manager. `Manager (MRM)`: The node acts as the Media Redundancy Manager. It monitors the ring and controls the blocking/unblocking of ring ports. `Auto Manager (MRA)`: The node participates in the manager selection protocol. If the current manager fails, an Auto Manager can take over the role
Name	Specifies the CFM domain configured for the ring. This identifies the control VLAN used for MRP Protocol Data Units (PDUs). The domain controls which VLAN tags the ring ports transmit.
Port1	Selects the port interface on the switch that represents MRP port1. If attempting to enable the MRP instance without an assigned Port1 interface, an error will be shown. Default is none, indicating that no port has been selected.
Port2	See description of Port1.
Priority	Sets the priority of the MRA/MRM. Used when operational role is MRM. The lower numerical value, the higher priority. As discussed in the last paragraph of the MRA section, the priority is really not relevant when the configured role is MRM. The standard, however, mandates the priority of nodes configured as MRMs to be 0x0000, 0x1000-0x7000, or 0x8000. And for nodes configured as MRAs it needs to be0x9000-0xF000 or 0xFFFF. The no-form sets it to 0xA000, assuming the configured role is MRA.
VLAN	By default, all MRP_Xxx PDUs are transmitted VLAN untagged (a value of 0 in this field) By setting this field to a non-zero value, a VLAN tag with the provided value can be inserted into the PDUs. The end user must ensure that all ring ports are members of this VLAN ID. The TPID of the VLAN tag follows the ring ports' TPID, and the PCP will always be 7.

	The standard mandates that reception of both tagged - with any VLAN ID - and untagged PDUs must be processed by the MRP implementation, so there is no check for actual VLAN ID. However, an operational warning may be issued if, for instance, a ring port is not a member of the given VLAN ID.
Recovery Profile	Sets the target recovery time for the ring. Options typically include `10`, `30`, `200` or `500` milliseconds. Note: Faster profiles (10 ms, 30 ms) may require hardware acceleration and might not be available on all switches. The unsupported profiles will not be available in the drop-down list. All nodes in the ring must share the same profile.
Manager Reacts on Link Change	If checked, instructs the node operating as an MRM to react immediately on MRP_LinkDown PDUs in the closed ring state by unblocking the secondary ring port and send MRP_TopologyChange PDUs out on both ring ports. If unchecked, it instructs it to wait until it finds its own MRP_Test PDUs to be missing. Default is the unchecked.
OUI Type	An OUI is included in MRP_Test PDUs if configured role is MRA and in MRP_TestMgrNAck and MRP_TestPropagate PDUs. By default, the OUI is the three most significant bytes of the switch's MAC address. The OUI is as such not used for anything but identification of switch manufacturer/MRP protocol developer. Default is `Default`. *Note: Due to a bug in Wireshark's dissector of MRP PDUs, which manifests itself in the fact that it only can dissect PDUs with the Siemens OUI (08-00-06), it is possible to change the OUI of these three PDU types to that of Siemens. And there is an option for changing to a custom OUI.*
OUI Value	When OUI Type is set to Custom, this defines the OUI to use in the PDUs mentioned under *OUI Type*. The format is six hexadecimal digits (e.g. `abcdef`).
UUID	Set the UUID/DomainID for this instance. The format is `xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx`, where x is a hexadecimal digit. The UUID/DomainID is used in MRP PDUs (encoded as 16 bytes) for identification of which ring a given PDU belongs to. It is recommended - but not required - that the DomainID/UUID is the same on all nodes in a given ring. Default is all Fs.

Ring Signal Fail Trigger

Signal Fail Trigger for Port1				Signal fail Trigger for Port2			
Type	Domain	Service	MEP ID	Type	Domain	Service	MEP ID
Link			0	Link			0

This table allows for configuring how the MRP Ring determines that a ring link has failed. It contains a section for Port1 and another section for Port2.

Parameter	Description
Type	Selects how to determine Signal Fail. Choose **`Link`** to use the port's link state to determine whether or not the connection to the link partner is up or down. Choose **`MEP`** to use a CFM MEP running on that port. If the specified MEP doesn't exist or for one or the other reason is not working, the MRP implementation will issue an operational warning and use the link state instead. The next three columns of the table are used to identify the MEP. Default is **`Link`**.
Domain	Only used if 'Type' is set to 'MEP'. Set the MEP's Domain Name in this field.
Service	Only used if 'Type' is set to 'MEP'. Set the MEP Domain's Service name in this field.
MEPID	Only used if 'Type' is set to 'MEP'. Set the MEP Domain's Service's MEP ID in this field.

Interconnection Configuration

Role	Name	Mode	Port	VLAN	Recovery Profile	IID
None		Link Check	none	0		0

Parameter	Description
Role	Configures the node's role in an interconnection topology (connecting two rings): `None`: The node is not part of an interconnection. `MIM` (Media Redundancy Interconnection Manager): The node manages the interconnection link. `MIC` (Media Redundancy Interconnection Client): The node acts as a client on the interconnection link The remainder of the columns in this table are only used if Role is Client or Manager.
Name	Set the name of this interconnection domain. This is only used for easy identification and has no effect on how the MRP instance operates. Default is an empty string.
Mode	As described under MIM above, the MIM or MIC may either operate in `Link Check` (default) or `Ring Check` mode.
Port	Selects the port interface on the switch that represents the interconnection port. Default is `none`, indicating that no port has been selected.
VLAN	By default, all MRP_InXxx PDUs are transmitted VLAN untagged (a value of `0` in this field) By setting this field to a non-zero value, a VLAN tag with the provided value can be inserted into the PDUs. The end user must ensure that both ring ports and the interconnection port are members of this VLAN ID. The TPID of the VLAN tag follows the egress port's TPID, and the PCP will always be 7. The standard mandates that reception of both tagged - with any VLAN ID - and untagged PDUs must be processed by the MRP implementation, so there is no check for actual VLAN ID. However, an operational warning may be issued if, for instance, one of the three ports is not a member of the given VLAN ID.
Recovery Profile	This selects the interconnection's recovery profile. Default is `500 ms`.
IID	The interconnection ID (IID) is a 16-bit number that uniquely identifies the four nodes making up this interconnection domain. All four nodes must use the same IID. Default is `0`.

Interconnection Signal Fail Trigger

Signal Fail Trigger for Interconnection Port			
Type	Domain	Service	MEPID
Link			0

This table allows for configuring how the interconnection ring determines that the interconnection port's link has failed. For a description of the parameters please refer to *Ring Signal Fail Trigger* of Port1 and Port2 as the fields are identical.

4.10 DHCPv4

DHCPv4 clients can request IP addresses automatically from a DHCP server, which then dynamically leases them for a defined period from an allocated IP address pool. This reduces manual configuration errors and simplifies address management, especially for large networks.

Received and transmitted DHCP packets are displayed under Monitor→DHCPv4→Detailed Statistics.
To enable dynamic IP address allocation for this switch, refer to Configuration→System→IP.

4.10.1 Server

The DHCPv4 Server is used to allocate network addresses and deliver configuration parameters to dynamically configured hosts called DHCP client.

4.10.1.1 Mode

The DHCPv4 Server can be enabled globally or per VLAN.
To monitor statistics for the DHCPv4 Server status can be viewed under, refer to Monitor→DHCPv4→Server.

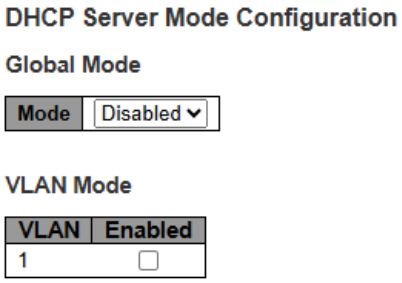

Global Mode enables or disables the DHCPv4 Server system wide via the drop-down menu.
VLAN Mode enables the DHCPv4 Server per VLAN by checking *Enabled* box according to the VLAN ID.

4.10.1.2 Excluded IP

The defined IP range will be excluded from the IP pool, and the DHCP Server will not allocate those to a DHCP client.
DHCPv4 Server status can be viewed under Monitor→DHCPv4→Server.

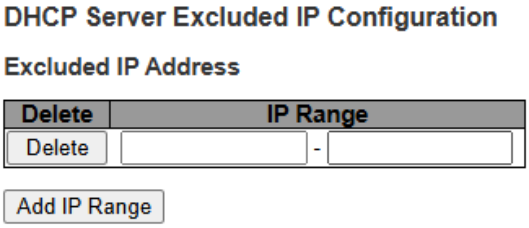

Defines the **IP Range** to be excluded IP addresses. The first excluded IP must be smaller than or equal to the second excluded IP. BUT, if the IP range contains only 1 excluded IP, then you can just input it to either one of the first and second excluded IP or both.

4.10.1.3 Pool

The DHCP pool is a range of IP addresses and associated network parameters that a DHCP server can assign to clients. It defines the address space, lease duration, and optional configuration settings such as subnet mask, gateway, and DNS servers.
DHCPv4 Server status can be viewed under Monitor→DHCPv4→Server.

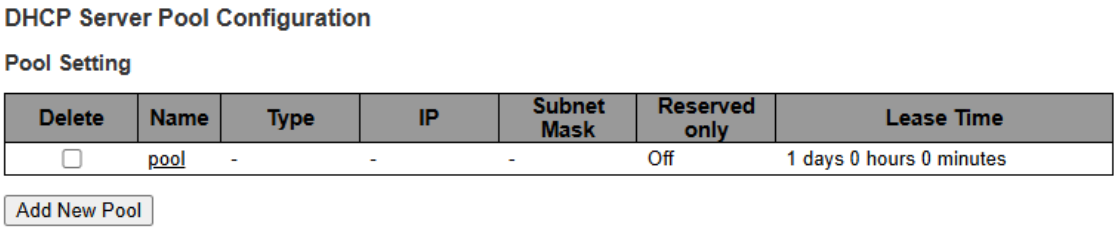

To configure a new DHCP IP pool follow these steps:
1. Click on *Add New Pool*
2. Assign a name for the pool
3. Click *Save*
4. Click on the name that you assigned
5. Configure the DHCP pool parameters and options

The IP pool will consist of all IP addresses within the subnet as specified by the *IP* and the *Subnet Mask*. This means that the size of the IP pool will always be a power of 2 (2, 4, 8, 16, 32, 64, 128, or 256). If the pool needs to be limited to a more specific number, this can be done at Configuration→DHCPv4→Server→Excluded IP.

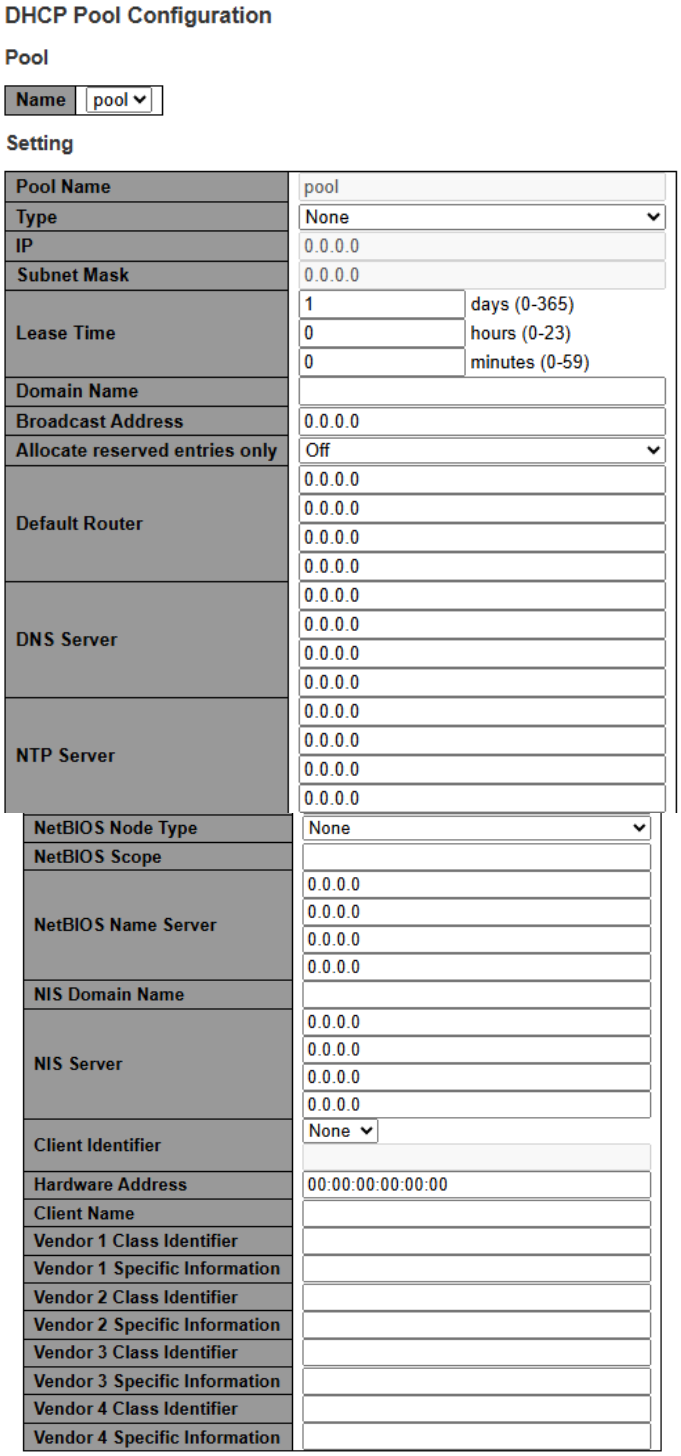

Parameter	Description
Pool Name	Displays the selected pool name.
Type	Specify which type of the pool is. **Network**: the pool defines a pool of IP addresses to service more than one DHCP client. **Host**: the pool services for a specific DHCP client identified by client identifier or hardware address.

IP	Specify the base IP of the DHCP address pool.
Subnet Mask	DHCP option 1. Specify subnet mask of the DHCP address pool.
Lease Time	DHCP option 51, 58 and 59. Specify lease time that allows the client to request a lease time for the IP address. If all are 0's, then it means the lease time is infinite.
Domain Name	DHCP option 15. Specify domain name that client should use when resolving hostname via DNS.
Broadcast Address	DHCP option 28. Specify the broadcast address in use on the client's subnet.
Allocate reserved entries only	Limits IP addresses obtainable from the pool to those entered into the reserved entries table. Select **On** to activate and **Off** to deactivate.
Default Router	DHCP option 3. Specify a list of IP addresses for routers on the client's subnet.
DNS Server	DHCP option 6. Specify a list of Domain Name System name servers available to the client.
NTP Server	DHCP option 42. Specify a list of IP addresses indicating NTP servers available to the client.
NetBIOS Node Type	DHCP option 46. Specify NetBIOS node type option to allow Netbios over TCP/IP clients which are configurable to be configured as described in RFC 1001/1002.
NetBIOS Scope	DHCP option 47. Specify the NetBIOS over TCP/IP scope parameter for the client as specified in RFC 1001/1002.
NetBIOS Name Server	DHCP option 44. Specify a list of NBNS name servers listed in order of preference.
NIS Domain Name	DHCP option 40. Specify the name of the client's NIS domain.
NIS Server	DHCP option 41. Specify a list of IP addresses indicating NIS servers available to the client.
Client Identifier	DHCP option 61. Specify client's unique identifier to be used when the pool is the type of host. Select the type of client identifier at first. **None**: client identifier is not specified yet. **Name**: the type of client identifier is other than hardware. **MAC**: the type of client identifier is MAC address.
Hardware Address	Specify client's hardware (MAC) address to be used when the pool is the type of host.
Client Name	DHCP option 12. Specify the name of client to be used when the pool is the type of host.
Vendor i Class Identifier	DHCP option 60. Specify to be used by DHCP client to optionally identify the vendor type and configuration of a DHCP client. DHCP server will deliver the corresponding option 43 specific information to the client that sends option 60 vendor class identifier.
Vendor i Specific Information	DHCP option 43. Specify vendor specific information according to option 60 vendor class identifier.

Individual IP adresses can be reserved to only be leased on a specific port. Refer to Monitor→Ports→Name Map for interface notation.

Reserved Ip Addresses

Delete	Reserved address	Interface
☐	192.168.12.1	Gi 1/1

Parameter	Description
Delete	Check this box to delete the entry when saving.

Reserved address	IP address to reserve for port interface.
Interface	Port interface attached to reserved IP address. This field can be updated.

4.10.2 Snooping

DHCP snooping is a security feature that protects the network by filtering untrusted DHCP messages and maintaining a DHCP snooping binding database (binding table). It acts as a firewall between untrusted hosts and DHCP servers by classifying interfaces as trusted (typically connected to DHCP servers or uplink switches) or untrusted (typically connected to end devices). DHCP snooping dynamically builds the binding table and blocks unauthorized DHCP packets, preventing clients from receiving IP addresses from rogue DHCP servers. The DHCP snooping binding database contains the MAC address, IP address, lease time, binding type, VLAN ID, and interface information associated with DHCP clients learned on the switch's untrusted interfaces.

When the switch receives a DHCP packet on an untrusted interface in a VLAN where DHCP snooping is enabled, it validates the packet. The switch can compare the source MAC address in the Ethernet header with the DHCP client hardware address (CHADDR field); if the addresses match (default behavior), the packet is forwarded, and if they do not match, the packet is dropped. In addition, DHCP responses (OFFER and ACK) received on untrusted interfaces are dropped to prevent rogue DHCP servers from assigning IP addresses to clients.

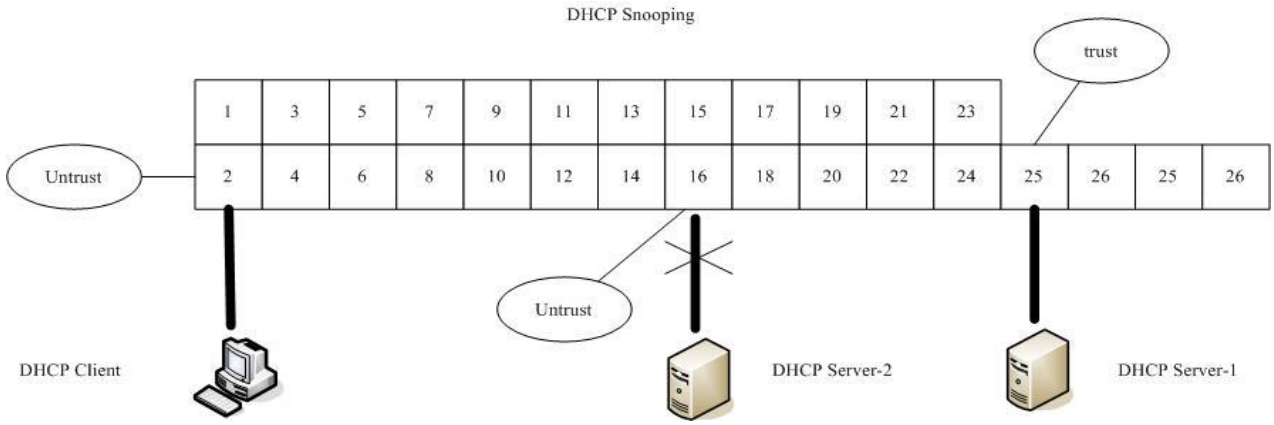

The switch drops a DHCP packet when one of these situations occurs:

- ✓ A DHCP server packet (ie: DHCPOFFER, DHCPACK, DHCPNAK, or DHCPLEASEQUERY) is received on an untrusted port.
- ✓ A packet is received on an untrusted interface, and the source MAC address and the DHCP client hardware address do not match any existing DHCP snooping binding entry.

Trusted ports are connected to DHCP servers or other switches that lead to a DHCP server. The switch discards DHCP packets from trusted ports only if the rate at which DHCP packets arrive is too high. The switch learns dynamic bindings from trusted ports.
Note: The switch will drop all DHCP requests if you enable DHCP snooping and there are no trusted ports.

Untrusted ports are connected to DHCP clients. The switch discards DHCP packets from un-trusted ports in the following situations:

- The packet is a DHCP server packet (for example, OFFER, ACK, or NACK).
- The source MAC address and source IP address in the packet do not match any existing DHCP snooping binding entry.
- The packet is a DHCPRELEASE or DHCPDECLINE packet, and the source MAC address and source port

do not match any existing binding entry.

- The rate at which DHCP packets arrive on the port exceeds the configured rate limit.

DHCP Snooping Database

The switch stores the binding table in volatile memory. If the switch restarts, it loads static bindings from permanent memory but loses the dynamic bindings, in which case the devices in the network have to send DHCP requests again. If the port link goes down, the entries learned by this port in the DHCP snooping binding table will be deleted.

The DHCP Snooping Table can be viewed at Monitor→DHCPv4→Snooping Table.

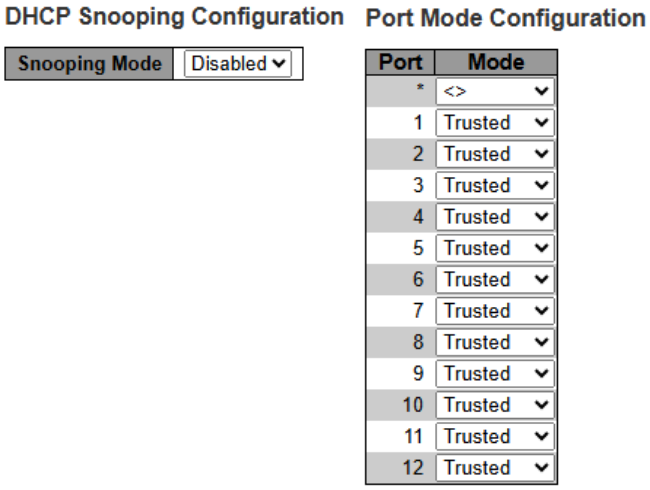

Parameter	Description
Snooping Mode	Set the DHCP snooping mode operation. Possible modes are: **Enabled**: Enable DHCP snooping mode operation. When DHCP snooping mode operation is enabled, the DHCP request messages will be forwarded to trusted ports and only allow reply packets from trusted ports. **Disabled**: Disable DHCP snooping mode operation.
Port Mode Configuration	Indicates the DHCP snooping port mode. Possible port modes are: **Trusted**: Configures the port as trusted source of the DHCP messages. **Untrusted**: Configures the port as untrusted source of the DHCP messages.

4.10.3 Relay

A DHCP relay agent is used to forward and to transfer DHCP messages between the clients and the server when they are not in the same subnet domain. The relay agent receives broadcast DHCP messages from clients and unicasts them to the configured server, enabling address assignment without a local DHCP server on each subnet. It stores the incoming interface IP address in the GIADDR field of the DHCP packet. The DHCP server can use the value of GIADDR field to determine the assigned subnet. For such condition, please make sure the switch configuration of VLAN interface IP address and PVID (Port VLAN ID) correctly.

- An IP address on the VLAN interface that serves as the default gateway for the DHCP clients in that VLAN.
- The PVID on access ports so that client ports are mapped to the correct VLAN.
- The DHCP Relay function on the corresponding VLAN interface, including the correct DHCP server IP address.

The DHCP Relay can be viewed at Monitor→DHCPv4→Relay Statistics.

DHCP Relay Configuration

Relay Mode	Disabled
Relay Server	0.0.0.0
Relay Information Mode	Disabled
Relay Information Policy	Keep

Parameter	Description
Relay Mode	Indicates the DHCP relay mode operation. Possible modes are: **Enabled**: Enable DHCP relay mode operation. When DHCP relay mode operation is enabled, the agent forwards and transfers DHCP messages between the clients and the server when they are not in the same subnet domain. And the DHCP broadcast message won't be flooded for security considerations. **Disabled**: Disable DHCP relay mode operation.
Relay Server	Indicates the DHCP relay server IP address.
Relay Information Mode	Indicates the DHCP relay information mode option operation. The option 82 circuit ID format as **[vlan_id][module_id][port_no]**. The first four characters represent the VLAN ID, the fifth and sixth characters are the module ID(in standalone device it always equal 0, in stackable device it means switch ID), and the last two characters are the port number. For example, **00030108** means the DHCP message receive form VLAN ID 3, switch ID 1, port No 8. And the option 82 remote ID value is equal the switch MAC address. Possible modes are: **Enabled**: Enable DHCP relay information mode operation. When DHCP relay information mode operation is enabled, the agent inserts specific information (option 82) into a DHCP message when forwarding to DHCP server and removes it from a DHCP message when transferring to DHCP client. It only works when DHCP relay operation mode is enabled. **Disabled**: Disable DHCP relay information mode operation.
Relay Information Policy	Indicates the DHCP relay information option policy. When DHCP relay information mode operation is enabled, if the agent receives a DHCP message that already contains relay agent information it will enforce the policy. The 'Replace' policy is invalid when relay information mode is disabled. Possible policies are: **Replace**: Replace the original relay information when a DHCP message that already contains it is received. **Keep**: Keep the original relay information when a DHCP message that already contains it is received. **Drop**: Drop the package when a DHCP message that already contains relay information is received.

4.11 DHCPv6

DHCPv6 is designed for IPv6 and differs from DHCPv4 in address assignment and identification. Clients are identified by a DUID rather than a MAC address, and DHCPv6 uses different message types and UDP ports (546/547) and relies on multicast rather than broadcast communication.

4.11.1 Snooping

For more information about DHCP Snooping refer to Configuration→DHCPv4→Snooping.
DHCPv6 Snooping status can be viewed under Monitor→DHCPv6→Snooping Table or Snooping Statistics.

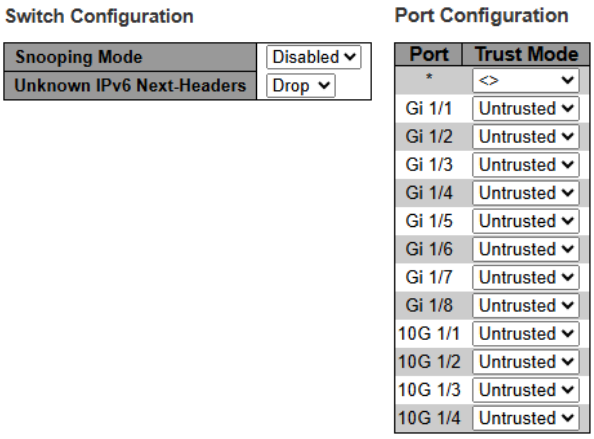

Parameter	Description
Snooping Mode	Set the DHCP snooping mode operation. Possible modes are: **Enabled**: Enable DHCPv6 snooping mode operation. When DHCPv6 snooping mode operation is enabled, the DHCP request messages will be forwarded to trusted ports and only allow reply packets from trusted ports. **Disabled**: Disable DHCPv6 snooping mode operation.
Unknown IPv6 Next-Headers	Set how Unknown IPv6 Next-Header values should be treated. The switch needs to parse all IPv6 packets to a DHCPv6 client to determine if it is in fact a DHCPv6 message. If an unknown IPv6 extension header is encountered the parsing cannot continue. See RFC 7610, section 5, item 3 for details. Possible options are: **Drop**: Drop packets with unknown IPv6 extension headers. This is the most secure option but may result in traffic disruptions. **Allow**: Allow packets with unknown IPv6 extension headers. This is a less secure option but prevents traffic disruptions.
Port Mode Configuration	Set the DHCPv6 snooping port mode. Possible port modes are: **Trusted**: Configures the port as trusted source of the DHCPv6 messages. **Untrusted**: Configures the port as untrusted source of the DHCPv6 messages.

4.11.2 Relay

DHCPv6 Relay status can be viewed at Monitor→DHCPv6→Relay.

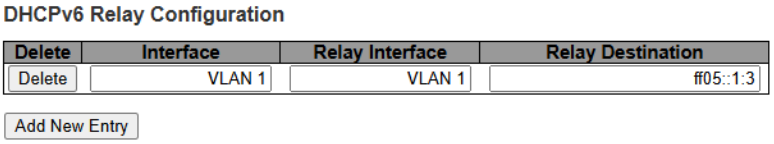

Parameter	Description
Interface	Interface identification.
Relay Interface	Interface identification. The id of the interface used for relaying.
Relay Destination	An Ipv6 address represented as human readable test as specified in RFC5952. The IPv6 address of the DHCPv6 server that requests shall be relayed to. The default value **ff05::1:3** means 'any DHCP server'.

4.12 Security

Layer 2 security features protect the network against threats that originate before IP-level controls can be applied. They prevent attacks such as MAC spoofing, ARP poisoning, rogue DHCP servers, and unauthorized

access at the switch port level. Implementing Layer 2 security is essential to ensure network integrity, device authentication, and reliable operation of higher-layer protocols.

4.12.1 Switch

Switch-level security mechanisms protect management access and operational data on the device. They control how users authenticate, what actions they are permitted to perform, and how management traffic is securely accessed and monitored. These features help prevent unauthorized configuration changes, protect sensitive information, and support auditing and diagnostics.

4.12.1.1 Users

This page provides an overview of the current users. Currently the only way to login as another user on the web server is to close and reopen the browser. Click on *Add New User* to create a new user or click on a user name to edit an existing user.

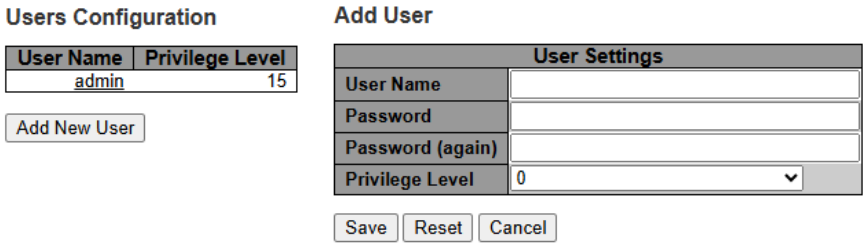

Parameter	Description
User Name	A string identifying the user name that this entry should belong to. The allowed string length is 1 to 31. The valid user name allows letters, numbers and underscores.
Password	The password of the user. The allowed string length is 0 to 31. Any printable characters including space is accepted.
Privilege Level	The privilege level of the user. The allowed range is **`0 to 15`**. If the privilege level value is 15, it can access all groups, i.e. that is granted the fully control of the device. But others value needs to refer to each group privilege level. User's privilege should be same or greater than the group privilege level to have the access of that group. By default setting, most groups privilege level 5 has the read-only access and privilege level 10 has the read-write access. And the system maintenance (software upload, factory defaults and etc.) need user privilege level 15. Generally, the privilege level 15 can be used for an administrator account, privilege level 10 for a standard user account and privilege level 5 for a guest account.

4.12.1.2 Privilege Levels

The privilege levels required for accessing configuration or monitoring information of features can be set per group. If a user's privilege level is equal to or higher than the required privilege level for a group and access mode, the access will be granted.

Privilege Level Configuration

Group Name	Privilege Levels			
	Configuration Read-only	Configuration/Execute Read/write	Status/Statistics Read-only	Status/Statistics Read/write
Aggregation	5	10	5	10
Alarm	5	10	5	10
APS	5	10	5	10
CFM	5	10	5	10
DDMI	5	10	5	10
Debug	15	15	15	15
DHCP	5	10	5	10
DHCPv6_Client	5	10	5	10
Diagnostics	5	10	5	10
ERPS	5	10	5	10
ETH_LINK_OAM	5	10	5	10
Firmware	5	10	5	10
Green_Ethernet	5	10	5	10
IP	5	10	5	10
IPMC_Snooping	5	10	5	10
LACP	5	10	5	10
LLDP	5	10	5	10
Loop_Protect	5	10	5	10
MAC_Table	5	10	5	10
Miscellaneous	15	15	15	15
MRP	5	10	5	10
MVR	5	10	5	10
NTP	5	10	5	10
POE	5	10	5	10
Ports	5	10	1	10
Private_VLANs	5	10	5	10
PTP	5	10	5	10
QoS	5	10	5	10
RMirror	5	10	5	10
Security(access)	10	10	5	10
Security(network)	5	10	5	10
sFlow	5	10	5	10
Spanning_Tree	5	10	5	10
System	5	10	1	10
UDLD	5	10	5	10
uFDMA_AIL	5	10	5	10
uFDMA_CIL	5	10	5	10
UPnP	5	10	5	10
VCL	5	10	5	10
VLAN_Translation	5	10	5	10
VLANs	5	10	5	10
Voice_VLAN	5	10	5	10
XXRP	5	10	5	10

Parameter	Description
Group Name	The name identifying the privilege group. In most cases, a privilege level group consists of a single module (e.g. LACP, RSTP or QoS), but a few of them contains more than one. The following description defines these privilege level groups in details: **System**: Contact, Name, Location, Time zone, Daylight Saving Time, Log. **Security (access)**: Users, Privilege Levels, Auth Method, SSH, HTTPS, Access Management, RADIUS, TACACS+. **Security (network)**: RMON, WEB, Port Security, NAS, ACL, IPv4/IPv6 Source Guard, ARP Inspection. **Diagnostics**: MIB Retrieval, VeriPHY. **Miscellaneous**: CLI- System Reboot, System Restore Default, System Password, Configuration Save, Configuration Load and Firmware Load. Web- Users, Privilege Levels and everything in Maintenance (except Image Select and Save startup-config). **Debug**: Only present in CLI.
Privilege Levels	The Privilege Levels can be configured between **0 to 15** (where 0 is lowest level and 15 is highest level) Every group has an authorization Privilege level for the following sub groups: configuration read-only, configuration/execute read-write, status/statistics read-only, status/statistics read-write (e.g. for clearing of statistics). User Privilege should be same or greater than the authorization Privilege level to have the access to that group.

4.12.1.3 Auth Method

The authentication section allows you to configure how a user is authenticated when he logs into the switch via one of the management client interfaces.

Remote authentication servers can be configured at Configuration→Security→AAA.

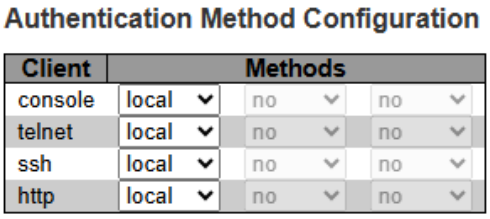

Parameter	Description
Client	The management client for which the configuration below applies.
Methods	Method can be set to one of the following values: **no**: Authentication is disabled and login is not possible. **local**: Use the local user database on the switch stack for authentication. **radius**: Use remote RADIUS server(s) for authentication. **tacacs**: Use remote TACACS+ server(s) for authentication. Methods that involve remote servers are timed out if the remote servers are offline. In this case the next method is tried. Each method is tried from left to right and continues until a method either approves or rejects a user. If a remote server is used for primary authentication it is recommended to configure secondary authentication as 'local'. This will enable the management client to login via the local user database if none of the configured authentication servers are alive.

The command authorization section allows you to limit the CLI commands available to a user depending on the CLI client used to access the switch.

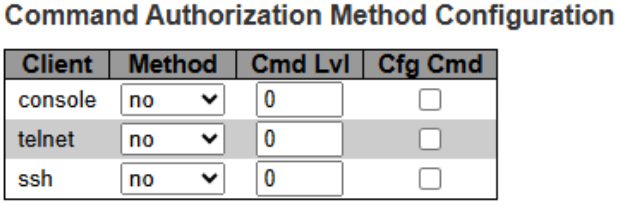

Parameter	Description
Client	The management client for which the configuration below applies.
Method	Method can be set to one of the following values: **no**: Command authorization is disabled. User is granted access to CLI commands according to his privilege level. **tacacs**: Use remote TACACS+ server(s) for command authorization. If all remote servers are offline, the user is granted access to CLI commands according to his privilege level.
Cmd Lvl	Authorize all commands with a privilege level higher than or equal to this level. Valid values are in the range **0 to 15**.
Cfg Cmd	Also authorize configuration commands. If disabled, configuration command will be authorized using the switch user privilege level.

The accounting section allows you to configure command and exec (login) accounting. If accounting is enabled, user commands are logged to a remote accounting server, providing data for auditing, monitoring, and troubleshooting.

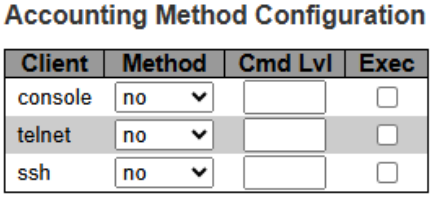

Parameter	Description
Client	The management client for which the configuration below applies.
Methods	Method can be set to one of the following values: **no**: Accounting is disabled. **tacacs**: Use remote TACACS+ server(s) for accounting.
Cmd Lvl	Enable accounting of all commands with a privilege level higher than or equal to this level. Valid values are in the range **0 to 15**. Leave the field empty to disable command accounting.
Exec	Enable exec (login) accounting.

4.12.1.4 SSH

The CLI access to the switch can be **enabled** or **disabled**.

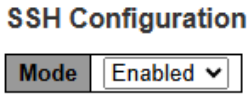

4.12.1.5 HTTPS

The HTTPS page controls secure Web GUI access and certificate maintenance.

Mode selects whether HTTPS service is Disabled or Enabled. Automatic Redirect selects whether HTTP access is redirected to HTTPS when the secure service is enabled.

Certificate Maintain supports None, Delete, Upload, and Generate operations. Upload installs a certificate selected from the browser or URL. Generate creates a new self-signed certificate on the switch. Delete removes the existing certificate and allows the switch to regenerate or use another configured certificate.

Certificate Status shows whether the switch secure HTTP certificate is present. When a self-signed certificate is used, the browser may display a security warning. Verify the device address and certificate details before continuing.

Figure: HTTPS configuration page showing Mode, Automatic Redirect, Certificate Maintain, and Certificate Status.

Parameter	Description

Mode	Enables or disables the HTTPS service for secure Web GUI access.
Automatic Redirect	Controls whether HTTP access is redirected to HTTPS when HTTPS service is enabled.
Certificate Maintain	Selects the certificate maintenance operation. None keeps the current state. Delete removes the installed certificate. Upload installs a selected certificate. Generate creates a new self-signed certificate on the switch.
Certificate Status	Shows whether the switch secure HTTP certificate is present. If a self-signed certificate is used, the browser may display a certificate warning.
Save	Applies the HTTPS configuration changes.
Reset	Discards unsaved local changes and restores the displayed settings to the current saved configuration.

4.12.1.6 Access Management

Management access can be allowed for only selected ranges of IP addresses. The maximum number of entries is **16**. If the access type matches any one of the access management entries, it will allow access to the switch.

For more Access Management Statistics refer to Monitor→Security→Switch→Access Management Statistics.

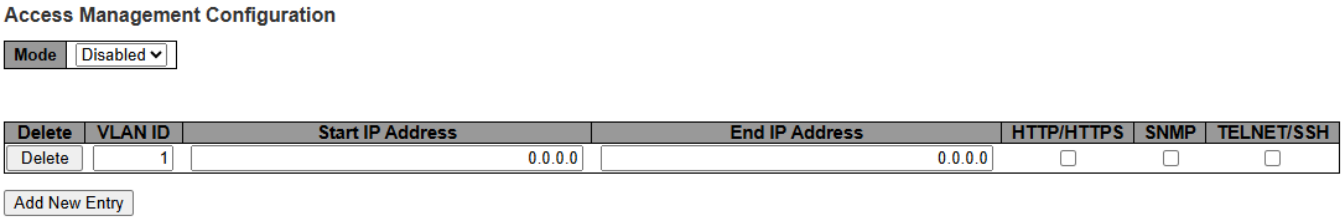

Parameter	Description
Mode	Indicate the access management mode operation. Possible modes are: **Enabled**: Enable access management mode operation. **Disabled**: Disable access management mode operation.
Delete	Check to delete the entry. It will be deleted during the next save.
VLAN ID	Indicates the VLAN ID for the access management entry.
Start IP address	Indicates the start IP unicast address for the access management entry.
End IP address	Indicates the end IP unicast address for the access management entry.
HTTP/HTTPS	Indicates that the host can access the switch from HTTP/HTTPS interface (Web GUI) if the host IP address matches the IP address range provided in the entry.
SNMP	Indicates that the host can access the switch from SNMP interface if the host IP address matches the IP address range provided in the entry.
TELNET/SSH	Indicates that the host can access the switch from TELNET/SSH interface if the host IP address matches the IP address range provided in the entry.

4.12.1.7 RMON

RMON is a standard for monitoring network traffic and performance on Ethernet devices. It collects statistical data, events, and alarms locally on the device and reports them to a network management system, reducing the need for continuous polling.

Statistics

Individual ports can be added here to have their RMON statistics displayed in the monitor section.
RMON Statistics can be viewed at Monitor→Security→Switch→RMON→Statistics.

RMON Statistics Configuration

Delete	ID	Data Source
Delete		.1.3.6.1.2.1.2.2.1.1. 0

Add New Entry | Save | Reset

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
ID	Indicates the index of the entry. The range is from **1 to 65535**.
Data Source	Indicates the port ID which wants to be monitored. If in stacking switch, the value must add 1000000*(switch ID) + Port No, for example, for port 5 on switch 1, the value is 1000005.

History

Individual ports can be added here to have their RMON history displayed in the monitor section.
RMON History can be viewed at Monitor→Security→Switch→RMON→History.

RMON History Configuration

Delete	ID	Data Source	Interval	Buckets	Buckets Granted
Delete		.1.3.6.1.2.1.2.2.1.1. 0	1800	50	

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
ID	Indicates the index of the entry. The range is from **1 to 65535**.
Data Source	Indicates the port ID which wants to be monitored. If in stacking switch, the value must add 1000000*(switch ID) + Port No, for example, if the port is switch 1 port 5, the value is 1000005.
Interval	Indicates the interval in seconds for sampling the history statistics data. The range is from **1 to 3600**, default value is 1800 seconds.
Buckets	Indicates the maximum data entries associated this History control entry stored in RMON. The range is from **1 to 50**, default value is 50.
Buckets Granted	The number of data shall be saved in the RMON.

Alarm

Individual ports can be added here to have their RMON alarms displayed in the monitor section.
RMON Alarm can be viewed at Monitor→Security→Switch→RMON→Alarm.

RMON Alarm Configuration

Delete	ID	Interval	Variable	Sample Type	Value	Startup Alarm	Rising Threshold	Rising Index	Falling Threshold	Falling Index

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
ID	Indicates the index of the entry. The range is from **1 to 65535**.
Interval	Indicates the interval in seconds for sampling and comparing the rising and falling threshold. The range is from 1 to 2^31-1.
Variable	Indicates the particular OID to be sampled. The parameter format is **X.Y**, with **X** representing the variable type and **Y** the port ID ((switch ID)+ 1000000 + Port No.). The possible variable types are:

	10 - InOctets: The total number of octets received on the interface, including framing characters. **11 - InUcastPkts:** The number of unicast packets delivered to a higher-layer protocol. **12 - InNUcastPkts**: The number of broad-cast and multi-cast packets delivered to a higher-layer protocol. **12 - InDiscards**: The number of inbound packets that are discarded even the packets are normal. **14 - InErrors**: The number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol. **15 - InUnknownProtos**: the number of the inbound packets that were discarded because of the unknown or un-support protocol. **16 - OutOctets**: The number of octets transmitted out of the interface , including framing characters. **17 - OutUcastPkts**: The number of unicast packets that request to transmit. **18 - OutNUcastPkts**: The number of broad-cast and multi-cast packets that request to transmit. **19 - OutDiscards**: The number of outbound packets that are discarded event the packets is normal. **20 - OutErrors**: The number of outbound packets that could not be transmitted because of errors. **21 - OutQLen**: The length of the output packet queue (in packets).
Sample Type	The method of sampling the selected variable and calculating the value to be compared against the thresholds, possible sample types are: **Absolute**: Get the sample directly. **Delta**: Calculate the difference between samples (default).
Value	The value of the statistic during the last sampling period.
Startup Alarm	The method of sampling the selected variable and calculating the value to be compared against the thresholds, possible sample types are: **Rising:** Trigger alarm when the first value is larger than the rising threshold. **Falling:** Trigger alarm when the first value is less than the falling threshold. **RisingOrFalling:** Trigger alarm when the first value is larger than the rising threshold or less than the falling threshold (default).
Rising Threshold	Rising threshold value (**-2147483648 to 2147483647**).
Rising Index	Rising event index (**0 to 65535**). If this value is zero, no associated event will be generated, as zero is not a valid event index.
Falling Threshold	Falling threshold value (**-2147483648 to 2147483647**)
Falling Index	Falling event index (**0 to 65535**). If this value is zero, no associated event will be generated, as zero is not a valid event index.

Event

Individual ports can be added here to have their RMON events displayed in the monitor section.
RMON Event can be viewed at Monitor→Security→Switch→RMON→Event.

RMON Event Configuration

Delete	ID	Desc	Type	Event Last Time
Delete			none	0

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
ID	Indicates the index of the entry. The range is from **1 to 65535**.
Desc	Indicates this event, the string length is from **0 to 127**, default is a null string.
Type	Indicates the notification of the event, the possible types are: **none**: No SNMP log is created, no SNMP trap is sent. **log**: Create SNMP log entry when the event is triggered.

	snmptrap: Send SNMP trap when the event is triggered. **logandtrap**: Create SNMP log entry and sent SNMP trap when the event is triggered.
Event Last Time	Indicates the value of *sysUpTime* at the time this event entry last generated an event.

4.12.1.8 Web

Specify the number of seconds until the user is automatically logged out of the Web GUI. After the specified period of inactivity, the cookie exchanged during the web session will expire, and the user needs to re-enter credentials again in order to access the Web GUI again. Specifying 0 causes the web session not to timeout. The allowed range is 0 to 3600 seconds. The default value is 1800 (30 minutes).

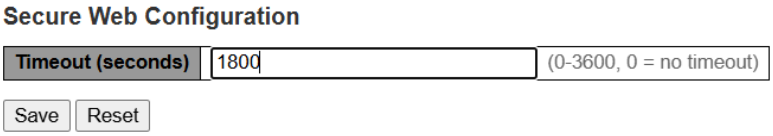

4.12.2 Network

This section provides features to control access and protect the network from unauthorized or malicious traffic. It includes port-based access controls, packet filtering, and IP/MAC validation mechanisms. These functions help prevent spoofing, unauthorized access, and common Layer 2 and Layer 3 attacks.

4.12.2.1 Port Security

Port Security restricts network access by validating MAC addresses learned on a switch port and limiting the number of devices that can connect to a port simultaneously. It prevents unauthorized devices from connecting and can take actions such as blocking traffic or shutting down the port when violations occur.

Configuration

Port Security allows for limiting the number of users on a given port. A user is identified by a MAC address and VLAN ID. If Port Security is enabled on a port, the Limit specifies the maximum number of users on the port. If this number is exceeded, an action is taken depending on the Violation Mode. The violation mode can be one of the four different described below.

Port Security Configuration can be viewed at Monitor→Security→Network→Port Security→Overview.

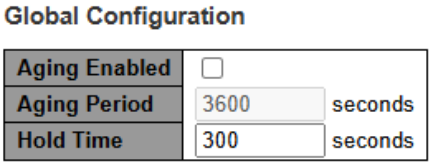

Parameter	Description
Aging Enabled	If checked, secured MAC addresses are subject to aging as discussed under *Aging Period*.
Aging Period	If *Aging Enabled* is checked, then the aging period is controlled with this input. If other modules are using the underlying functionality for securing MAC addresses, they may have other requirements to the aging period. The underlying functionality will use the shorter requested aging period of all modules that have aging enabled. The Aging Period can be set to a number between **10 and 10000000** seconds with a default of 3600 seconds. To understand why aging may be desired, consider the following scenario: Suppose an end-host is connected to a 3rd party switch or hub, which in turn is connected to a port on this switch on which Port Security is enabled. The end-host will be allowed to forward if the limit is not exceeded. Now suppose that the end-host logs off or powers down. If it wasn't for aging, the end-host would still

	take up resources on this switch and will be allowed to forward. To overcome this situation, enable aging. With aging enabled, a timer is started once the end-host gets secured. When the timer expires, the switch starts looking for frames from the end-host, and if such frames are not seen within the next Aging Period, the end-host is assumed to be disconnected, and the corresponding resources are freed on the switch.
Hold Time	The hold time - measured in seconds - is used to determine how long a MAC address is held in the MAC table if it has been found to violate the limit. Valid range is between **10 and 10000000** seconds with a default of 300 seconds. The reason for holding a violating MAC address in the MAC table is primarily to ensure that the same MAC address doesn't give rise to continuous notifications (if notifications on violation count is enabled).

Port Configuration

Port	Mode	Limit	Violation Mode	Violation Limit	Sticky	State
*	<>	4	<>	4		
1	Disabled	4	Protect	4		Disabled
2	Disabled	4	Protect	4		Disabled
3	Disabled	4	Protect	4		Disabled
4	Disabled	4	Protect	4		Disabled
5	Disabled	4	Protect	4		Disabled
6	Disabled	4	Protect	4		Disabled
7	Disabled	4	Protect	4		Disabled
8	Disabled	4	Protect	4		Disabled
9	Disabled	4	Protect	4		Disabled
10	Disabled	4	Protect	4		Disabled
11	Disabled	4	Protect	4		Disabled
12	Disabled	4	Protect	4		Disabled

Parameter	Description
Port	The port number to which the configuration below applies.
Mode	Controls whether Port Security is enabled on this port. Notice that other modules may still use the underlying port security features without enabling Port Security on a given port.
Limit	The maximum number of MAC addresses that can be secured on this port. This number cannot exceed 1023. Default is 4. If the limit is exceeded, an action is taken corresponding to the *violation mode*. The switch is "born" with a total number of MAC addresses from which all ports draw whenever a new MAC address is seen on a Port Security-enabled port. Since all ports draw from the same pool, it may happen that a configured maximum cannot be granted if the remaining ports have already used all available MAC addresses.
Violation Mode	If *Limit* is reached, the switch can take one of the following actions: **Protect**: Do not allow more than *Limit* MAC addresses on the port but take no further action. **Restrict**: If *Limit* is reached, subsequent MAC addresses on the port will be counted and marked as violating. Such MAC addresses are removed from the MAC table when the *hold time* expires. At most *Violation Limit* MAC addresses can be marked as violating at any given time. **Shutdown**: If *Limit* is reached, one additional MAC address will cause the port to be shut down. This implies that all secured MAC addresses be removed from the port, and no new addresses be learned. There are three ways to re-open the port: 1) In the Configuration→Ports page's "Configured" column, first disable the port, then restore the original mode. 2) Make a Port Security configuration change on the port. 3) Boot the switch.
Violation Limit	The maximum number of MAC addresses that can be marked as violating on this port. This number cannot exceed 1023. Default is 4. It is only used when *Violation Mode* is **Restrict**.
Sticky	Enables sticky learning of MAC addresses on this port. When the port is in sticky mode, all MAC addresses that would otherwise have been learned as dynamic are learned as sticky. Sticky MAC addresses are part of the running-config and can therefore be saved to startup-config. Sticky MAC addresses survive link changes (in contrast to Dynamic, which will have to be learned again). They also survive reboots if running-config is saved to startup-config. A port can be Sticky-enabled whether or not Port Security is enabled on that interface. In that way, it is possible to add sticky MAC addresses managementwise before enabling Port Security. To do that, use the Configuration→Security→Network→Port Security→MAC Addresses page.

MAC Address

On this page, you may add and delete static and sticky MAC addresses managed by Port Security.
Port Security MAC Address can be viewed at Monitor→Security→Network→Port Security→Details.

Port security defines three types of MAC addresses, of which static and sticky can be added or removed on this page:

- **Dynamic**: A MAC address learned through learn frames coming to the Port Security module while the interface in question is not in sticky mode. Dynamic entries disappear if they age out or if the interface link goes down.
- **Static**: A MAC address added by end-user through management. Static MAC addresses are not subject to aging and will be added to the MAC address table once Port Security gets enabled on the interface. Static entries are part of the running-config and will survive interface link state changes and reboots if saved to startup-config. Static entries can be added to the running-config at any time whether or not Port Security is enabled.
- **Sticky**: When the interface is in sticky mode, all entries that would otherwise have been learned as dynamic are learned as sticky.
  Like static entries, sticky entries are part of the running-config and will survive interface link state changes and reboots if saved to the startup-config.
  Though not the intention with Sticky entries, they can be added by management to the running-config at any time whether or not Port Security is enabled on the interface, as long as the interface is in Sticky mode. Sticky entries will disappear if the interface is taken out of Sticky mode.

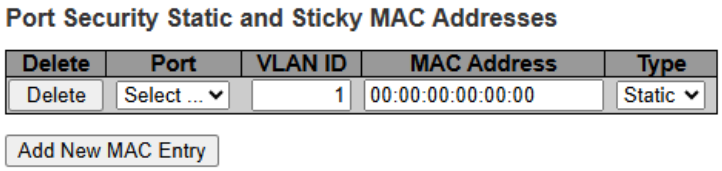

Parameter	Description
Delete	Press this button to remove the entry from the MAC address table (if present) and the running-config. Notice that dynamic entries may be removed all-together on an interface through Monitor→Security→Port Security→Switch and one-by-one through Monitor→Security→Port Security→Port
Port	The port number to which this MAC address is bound.
VLAN ID & MAC Address	The VLAN ID and MAC address in question.
Type	Indicates the type of entry and may be either Static or Sticky (see description above).

4.12.2.2 NAS

Network Access Server (NAS) allows you to configure the IEEE 802.1X and MAC-based authentication system and port settings. The IEEE 802.1X standard defines a port-based access control procedure that prevents unauthorized access to a network by requiring users to first submit credentials for authentication. One or more central servers, determine whether the user is allowed access to the network. These backend (RADIUS) servers are configured on the Configuration→Security→AAA page. The IEEE802.1X standard defines port-based operation, but non-standard variants overcome security limitations as shall be explored below.

MAC-based authentication allows for authentication of more than one user on the same port and doesn't require the user to have special 802.1X supplicant software installed on his system. The switch uses the user's

MAC address to authenticate against the backend server. Intruders can create counterfeit MAC addresses, which makes MAC-based authentication less secure than 802.1X authentication.

NAS can be viewed at Monitor→Security→Network→NAS.

System Configuration

Mode	Disabled
Reauthentication Enabled	☐
Reauthentication Period	3600 seconds
EAPOL Timeout	30 seconds
Aging Period	300 seconds
Hold Time	10 seconds
RADIUS-Assigned QoS Enabled	☐
RADIUS-Assigned VLAN Enabled	☐
Guest VLAN Enabled	☐
Guest VLAN ID	1
Max. Reauth. Count	2
Allow Guest VLAN if EAPOL Seen	☐

Parameter	Description
Mode	Indicates if NAS is globally enabled or disabled on the switch. If globally disabled, all ports are allowed forwarding of frames.
Reauthentication Enabled	If checked, successfully authenticated supplicants/clients are reauthenticated after the interval specified by the Reauthentication Period. Reauthentication for 802.1X-enabled ports can be used to detect if a new device is plugged into a switch port or if a supplicant is no longer attached. For MAC-based ports, reauthentication is only useful if the RADIUS server configuration has changed. It does not involve communication between the switch and the client, and therefore doesn't imply that a client is still present on a port (see *Aging Period* below).
Reauthentication Period	Determines the period, in seconds, after which a connected client must be reauthenticated. This is only active if the Reauthentication Enabled checkbox is checked. Valid values are in the range **`1 to 3600`** seconds.
EAPOL Timeout	Determines the time for retransmission of Request Identity EAPOL frames. Valid values are in the range **`1 to 65535`** seconds. This has no effect for MAC-based ports.
Aging Period	This setting applies to the following modes, i.e. modes using the Port Security functionality to secure MAC addresses: • Single 802.1X • Multi 802.1X • MAC-Based Auth. When the NAS module uses the Port Security module to secure MAC addresses, the Port Security module needs to check for activity on the MAC address in question at regular intervals and free resources if no activity is seen within a given period of time. This parameter controls exactly this period and can be set to a number between **`10 and 1000000`** seconds. If *reauthentication* is enabled and the port is in an 802.1X-based mode, this is not so critical, since supplicants that are no longer attached to the port will get removed upon the next reauthentication, which will fail. But if reauthentication is not enabled, the only way to free resources is by aging the entries. For ports in MAC-based Auth. mode, *reauthentication* doesn't cause direct communication between the switch and the client, so this will not detect whether the client is still attached or not, and the only way to free any resources is to age the entry.
Hold Time	This setting applies to the following modes, i.e. modes using the Port Security functionality to secure MAC addresses: • Single 802.1X • Multi 802.1X • MAC-Based Auth. If a client is denied access - either because the RADIUS server denies the client access or because the RADIUS server request times out (according to the timeout specified on the Configuration→Security→AAA page) - the client is put on hold in the Unauthorized state. The hold timer does not count during an on-going authentication.

	In MAC-based Auth. mode, the switch will ignore new frames coming from the client during the hold time. The Hold Time can be set to a number between **10 and 1000000** seconds.
RADIUS-Assigned QoS Enabled	RADIUS-assigned QoS provides a means to centrally control the traffic class to which traffic coming from a successfully authenticated supplicant is assigned on the switch. The RADIUS server must be configured to transmit special RADIUS attributes to take advantage of this feature (see RADIUS-Assigned QoS Enabled below for a detailed description). The "*RADIUS-Assigned QoS Enabled*" checkbox provides a quick way to globally enable/disable RADIUS-server assigned QoS Class functionality. When checked, the individual ports' ditto setting determines whether RADIUS-assigned QoS Class is enabled on that port. When unchecked, RADIUS-server assigned QoS Class is disabled on all ports.
RADIUS-Assigned VLAN Enabled	RADIUS-assigned VLAN provides a means to centrally control the VLAN on which a successfully authenticated supplicant is placed on the switch. Incoming traffic will be classified to and switched on the RADIUS-assigned VLAN. The RADIUS server must be configured to transmit special RADIUS attributes to take advantage of this feature (see RADIUS-Assigned VLAN Enabled below for a detailed description). The "*RADIUS-Assigned VLAN Enabled*" checkbox provides a quick way to globally enable/disable RADIUS-server assigned VLAN functionality. When checked, the individual ports' ditto setting determines whether RADIUS-assigned VLAN is enabled on that port. When unchecked, RADIUS-server assigned VLAN is disabled on all ports.
Guest VLAN Enabled	A Guest VLAN is a special VLAN - typically with limited network access - on which 802.1X-unaware clients are placed after a network administrator-defined timeout. The switch follows a set of rules for entering and leaving the Guest VLAN as listed below. The "Guest VLAN Enabled" checkbox provides a quick way to globally enable/disable Guest VLAN functionality. When checked, the individual ports' ditto setting determines whether the port can be moved into Guest VLAN. When unchecked, the ability to move to the Guest VLAN is disabled on all ports.
Guest VLAN ID	This is the value that a port's Port VLAN ID is set to if a port is moved into the Guest VLAN. It is only changeable if the Guest VLAN option is globally enabled. Valid values are in the range [1; 4095].
Max. Reauth. Count	The number of times the switch transmits an EAPOL Request Identity frame without response before considering entering the Guest VLAN is adjusted with this setting. The value can only be changed if the Guest VLAN option is globally enabled. Valid values are in the range [1; 255].
Allow Guest VLAN if EAPOL Seen	The switch remembers if an EAPOL frame has been received on the port for the life-time of the port. Once the switch considers whether to enter the Guest VLAN, it will first check if this option is enabled or disabled. If disabled (unchecked; default), the switch will only enter the Guest VLAN if an EAPOL frame has not been received on the port for the life-time of the port. If enabled (checked), the switch will consider entering the Guest VLAN even if an EAPOL frame has been received on the port for the life-time of the port. The value can only be changed if the Guest VLAN option is globally enabled.

Port Configuration

The following authorization modes are available to select as *Admin State*:

Force Authorized

Sends one EAPOL Success frame when the port links up. All clients gain network access without authentication.

Force Unauthorized

Sends one EAPOL Failure frame when port link-up occurs, denying network access to all clients on this port.

Port-based 802.1X

In 802.1X, the end device (Supplicant) authenticates via the switch (Authenticator) to a RADIUS server (Authentication Server). The switch forwards EAPOL frames (from supplicant) and RADIUS packets (to server) without interpreting authentication methods (MD5, PEAP, TLS).

Authentication Flow:

1. Supplicant → Switch: EAPOL frames

2. Switch → RADIUS: Encapsulated EAP + port/IP attributes
3. RADIUS → Switch: Accept/Reject → Port opens/blocks

Note Server timeout must be shorter than supplicant's EAPOL Start retransmission interval to avoid authentication loops when primary RADIUS server is unresponsive.

Single 802.1X

In port-based 802.1X, successful authentication opens the entire port, allowing unauthenticated devices (e.g., via hub) to gain access—posing a security risk.

Single 802.1X resolves this by limiting authentication to one supplicant per port. The first supplicant attempting connection post-link-up gets authenticated via standard EAPOL. Only its MAC address (secured by Port Security) gains access, blocking others.

Multi 802.1X

Multi 802.1X is - like Single 802.1X - not an IEEE standard, but a variant that features many of the same characteristics. In Multi 802.1X, one or more supplicants can get authenticated on the same port at the same time. Each supplicant is authenticated individually and secured in the MAC table using the Port Security module.

In Multi 802.1X it is not possible to use the multicast BPDU MAC address as destination MAC address for EAPOL frames sent from the switch towards the supplicant, since that would cause all supplicants attached to the port to reply to requests sent from the switch. Instead, the switch uses the supplicant's MAC address, which is obtained from the first EAPOL Start or EAPOL Response Identity frame sent by the supplicant. An exception to this is when no supplicants are attached. In this case, the switch sends EAPOL Request Identity frames using the BPDU multicast MAC address as destination - to wake up any supplicants that might be on the port.

The maximum number of supplicants that can be attached to a port can be limited using the Port Security Limit Control functionality.

MAC-based Auth.

Unlike port-based 802.1X, MAC-based authentication is not a standard, but merely a best-practices method adopted by the industry. In MAC-based authentication, users are called clients, and the switch acts as the supplicant on behalf of clients. The initial frame (any kind of frame) sent by a client is snooped by the switch, which in turn uses the client's MAC address as both username and password in the subsequent EAP exchange with the RADIUS server. The 6-byte MAC address is converted to a string on the following form "**xx-xx-xx-xx-xx-xx**", that is, a dash (-) is used as separator between the lower-cased hexadecimal digits. The switch only supports the MD5-Challenge authentication method, so the RADIUS server must be configured accordingly.

When authentication is complete, the RADIUS server sends a success or failure indication, which in turn causes the switch to open up or block traffic for that particular client, using the Port Security module. Only then will frames from the client be forwarded on the switch. There are no EAPOL frames involved in this authentication, and therefore, MAC-based Authentication has nothing to do with the 802.1X standard.

The advantage of MAC-based authentication over 802.1X-based authentication is that the clients don't need special supplicant software to authenticate. The disadvantage is that MAC addresses can be spoofed by malicious users - equipment whose MAC address is a valid RADIUS user can be used by anyone. Also, only the MD5-Challenge method is supported. The maximum number of clients that can be attached to a port can be limited using the Port Security Limit Control functionality.

Port	Admin State	RADIUS-Assigned QoS Enabled	RADIUS-Assigned VLAN Enabled	Guest VLAN Enabled	Port State	Restart	
*	<>	☐	☐	☐			
1	Force Authorized	☐	☐	☐	Globally Disabled	Reauthenticate	Reinitialize
2	Force Authorized	☐	☐	☐	Globally Disabled	Reauthenticate	Reinitialize
3	Force Authorized	☐	☐	☐	Globally Disabled	Reauthenticate	Reinitialize
4	Force Authorized	☐	☐	☐	Globally Disabled	Reauthenticate	Reinitialize
5	Force Authorized	☐	☐	☐	Globally Disabled	Reauthenticate	Reinitialize
6	Force Authorized	☐	☐	☐	Globally Disabled	Reauthenticate	Reinitialize
7	Force Authorized	☐	☐	☐	Globally Disabled	Reauthenticate	Reinitialize
8	Force Authorized	☐	☐	☐	Globally Disabled	Reauthenticate	Reinitialize
9	Force Authorized	☐	☐	☐	Globally Disabled	Reauthenticate	Reinitialize
10	Force Authorized	☐	☐	☐	Globally Disabled	Reauthenticate	Reinitialize
11	Force Authorized	☐	☐	☐	Globally Disabled	Reauthenticate	Reinitialize
12	Force Authorized	☐	☐	☐	Globally Disabled	Reauthenticate	Reinitialize

Parameter	Description
Port	The port number for which the configuration below applies.
Admin State	If NAS is globally enabled, this selection controls the port's authentication mode. The following modes are available: **Force Authorized** **Force Unauthorized** **Port-based 802.1X** **Single 802.1X** **Multi 802.1X** **MAC-based Auth.** For more details, refer to the description above.
RADIUS-Assigned QoS Enabled	When RADIUS-Assigned QoS is both globally enabled and enabled (checked) on a given port, the switch reacts to QoS Class information carried in the RADIUS Access-Accept packet transmitted by the RADIUS server when a supplicant is successfully authenticated. If present and valid, traffic received on the supplicant's port will be classified to the given QoS Class. If (re-)authentication fails or the RADIUS Access-Accept packet no longer carries a QoS Class or it's invalid, or the supplicant is otherwise no longer present on the port, the port's QoS Class is immediately reverted to the original QoS Class (which may be changed by the administrator in the meanwhile without affecting the RADIUS-assigned). This option is only available for single-client modes, i.e. • Port-based 802.1X • Single 802.1X **RADIUS attributes used in identifying a QoS Class:** The User-Priority-Table attribute defined in RFC4675 forms the basis for identifying the QoS Class in an Access-Accept packet. Only the first occurrence of the attribute in the packet will be considered, and to be valid, it must follow this rule: • All 8 octets in the attribute's value must be identical and consist of ASCII characters in the range '0' - '7', which translates into the desired QoS Class in the range [0; 7].
RADIUS-Assigned VLAN Enabled	When RADIUS-Assigned VLAN is both globally enabled and enabled (checked) for a given port, the switch reacts to VLAN ID information carried in the RADIUS Access-Accept packet transmitted by the RADIUS server when a supplicant is successfully authenticated. If present and valid, the port's Port VLAN ID will be changed to this VLAN ID, the port will be set to be a member of that VLAN ID, and the port will be forced into VLAN unaware mode. Once assigned, all traffic arriving on the port will be classified and switched on the RADIUS-assigned VLAN ID. If (re-)authentication fails or the RADIUS Access-Accept packet no longer carries a VLAN ID or it's invalid, or the supplicant is otherwise no longer present on the port, the port's VLAN ID is immediately reverted to the original VLAN ID (which may be changed by the administrator in the meanwhile without affecting the RADIUS-assigned). This option is only available for single-client modes, i.e. • Port-based 802.1X • Single 802.1X For trouble-shooting VLAN assignments, use the Monitor→VLAN pages. These pages show which modules have (temporarily) overridden the current Port VLAN configuration.

	RADIUS attributes used in identifying a VLAN ID: RFC2868 and RFC3580 form the basis for the attributes used in identifying a VLAN ID in an Access-Accept packet. The following criteria are used: • The Tunnel-Medium-Type, Tunnel-Type, and Tunnel-Private-Group-ID attributes must all be present at least once in the Access-Accept packet. • The switch looks for the first set of these attributes that have the same Tag value and fulfil the following requirements (if Tag == 0 is used, the Tunnel-Private-Group-ID does not need to include a Tag): - Value of Tunnel-Medium-Type must be set to "IEEE-802" (ordinal 6). - Value of Tunnel-Type must be set to "VLAN" (ordinal 13). - Value of Tunnel-Private-Group-ID must be a string of ASCII chars in the range '0' - '9', which is interpreted as a decimal string representing the VLAN ID. Leading '0's are discarded. The final value must be in the range [1; 4095].
Guest VLAN Enabled	**Guest VLAN Operation:** When a Guest VLAN enabled port's link comes up, the switch starts transmitting EAPOL Request Identity frames. If the number of transmissions of such frames exceeds *Max. Reauth. Count* and no EAPOL frames have been received in the meanwhile, the switch considers entering the Guest VLAN. The interval between transmission of EAPOL Request Identity frames is configured with *EAPOL Timeout*. If *Allow Guest VLAN if EAPOL Seen* is enabled, the port will now be placed in the Guest VLAN. If disabled, the switch will first check its history to see if an EAPOL frame has previously been received on the port (this history is cleared if the port link goes down or the port's *Admin State* is changed), and if not, the port will be placed in the Guest VLAN. Otherwise, it will not move to the Guest VLAN, but continue transmitting EAPOL Request Identity frames at the rate given by EAPOL Timeout. Once in the Guest VLAN, the port is considered authenticated, and all attached clients on the port are allowed access on this VLAN. The switch will not transmit an EAPOL Success frame when entering the Guest VLAN. While in the Guest VLAN, the switch monitors the link for EAPOL frames, and if one such frame is received, the switch immediately takes the port out of the Guest VLAN and starts authenticating the supplicant according to the port mode. If an EAPOL frame is received, the port will never be able to go back into the Guest VLAN if the "Allow Guest VLAN if EAPOL Seen" is disabled.
Port State	The current state of the port. It can undertake one of the following values: **`Globally Disabled`:** NAS is globally disabled. **`Link Down`:** NAS is globally enabled, but there is no link on the port. **`Authorized`:** The port is in *Force Authorized* or a single-supplicant mode and the supplicant is authorized. **`Unauthorized`:** The port is in *Force Unauthorized* or a single-supplicant mode and the supplicant is not successfully authorized by the RADIUS server. **`X Auth/Y Unauth`:** The port is in a multi-supplicant mode. Currently X clients are authorized and Y are unauthorized.
Restart	Two buttons are available for each row. The buttons are only enabled when authentication is globally enabled and the port's *Admin State* is in an *EAPOL-based* or *MAC-based* mode. Clicking these buttons will not cause settings changed on the page to take effect. **`Reauthenticate`:** Schedules a reauthentication whenever the quiet-period of the port runs out (EAPOL-based authentication). For MAC-based authentication, reauthentication will be attempted immediately. The button only has effect for successfully authenticated clients on the port and will not cause the clients to get temporarily unauthorized. **`Reinitialize`:** Forces a reinitialization of the clients on the port and thereby a reauthentication immediately. The clients will transfer to the unauthorized state while the reauthentication is in progress.

4.12.2.3 ACL

ACL is an acronym for Access Control List. It is the list table of ACEs, containing access control entries that specify individual users or groups permitted or denied specific traffic objects, such as a process or a program. Each accessible traffic object contains an identifier to its ACL. The privileges determine whether there are specific traffic object access rights.

ACL implementations can be quite complex, for example, when the ACEs are prioritized for various situations. In networking, the ACL refers to a list of service ports or network services that are available on a host or server,

each with a list of hosts or servers permitted or denied using the service. ACL can generally be configured to control inbound traffic, and in this context, they are similar to firewalls.

There are 3 web-pages associated with the manual ACL configuration:
ACL → Access Control List: The web page shows the ACEs in a prioritized way, highest (top) to lowest (bottom). Default, the table is empty. An ingress frame will only get a hit on one ACE even though there are more matching ACEs. The first matching ACE will act (permit/deny) on that frame, and a counter associated with that ACE is incremented. An ACE can be associated with a Policy, 1 ingress port, or any ingress port (the whole switch). If an ACE Policy is created, then that Policy can be associated with a group of ports under the "Ports" web-page. There are number of parameters that can be configured with an ACE. Read the Web page help text to get further information for each of them. The maximum number of ACEs is 64.
ACL → Ports: The ACL Ports configuration is used to assign a Policy ID to an ingress port. This is useful to group ports to obey the same traffic rules. Traffic Policy is created under the "Access Control List" - page. You can you also set up specific traffic properties (Action / Rate Limiter / Port copy, etc) for each ingress port. They will though only apply if the frame gets past the ACE matching without getting matched. In that case a counter associated with that port is incremented. See the Web page to help text for each specific port property.
ACL → Rate Limiters: Under this page you can configure the rate limiters. There can be 15 different rate limiters, each ranging from 1-1024K packets per second. Under "Ports" and "Access Control List" webpages you can assign a Rate Limiter ID to the ACE(s) or ingress port(s).

Ports

Configure the ACL parameters (ACE) of each switch port. These parameters will affect frames received on a port unless the frame matches a specific ACE.
ACE is an acronym for **A**ccess **C**ontrol **E**ntry. It describes access permission associated with a particular ACE ID. There are three ACE frame types (Ethernet Type, ARP, and IPv4) and two ACE actions (permit and deny). The ACE also contains many detailed, different parameter options that are available for individual applications.
ACL status can be viewed at Monitor→Security→Network→ACL Status.

ACL Ports Configuration

Port	Policy ID	Action	Rate Limiter ID	Port Redirect	Mirror	Logging	Shutdown	State	Counter
*	0	<>	<>	Disabled Port 1 Port 2	<>	<>	<>	<>	*
1	0	Permit	Disabled	Disabled Port 1 Port 2	Disabled	Disabled	Disabled	Enabled	0
2	0	Permit	Disabled	Disabled Port 1 Port 2	Disabled	Disabled	Disabled	Enabled	0
3	0	Permit	Disabled	Disabled Port 1 Port 2	Disabled	Disabled	Disabled	Enabled	0
4	0	Permit	Disabled	Disabled Port 1 Port 2	Disabled	Disabled	Disabled	Enabled	0
5	0	Permit	Disabled	Disabled Port 1 Port 2	Disabled	Disabled	Disabled	Enabled	0
6	0	Permit	Disabled	Disabled Port 1 Port 2	Disabled	Disabled	Disabled	Enabled	0
7	0	Permit	Disabled	Disabled Port 1 Port 2	Disabled	Disabled	Disabled	Enabled	0
8	0	Permit	Disabled	Disabled Port 1 Port 2	Disabled	Disabled	Disabled	Enabled	39205
9	0	Permit	Disabled	Disabled Port 1 Port 2	Disabled	Disabled	Disabled	Enabled	0
10	0	Permit	Disabled	Disabled Port 1 Port 2	Disabled	Disabled	Disabled	Enabled	0
11	0	Permit	Disabled	Disabled Port 1 Port 2	Disabled	Disabled	Disabled	Enabled	0
12	0	Permit	Disabled	Disabled Port 1 Port 2	Disabled	Disabled	Disabled	Enabled	0

Parameter	Description
Port	The logical port for the settings contained in the same row.
Policy ID	Select the policy to apply to this port. The allowed values are **0** through **7**. The default value is 0.
Action	Select whether forwarding is permitted (**Permit**) or denied (**Deny**). The default value is **Permit**.
Rate Limiter ID	Select which rate limiter to apply on this port. The allowed values are Disabled or the values **1** through **15**. The default value is **Disabled**".
Mirror	Specify the mirror operation of this port. The allowed values are: **Enabled**: Frames received on the port are mirrored. **Disabled**: Frames received on the port are not mirrored. The default value is **Disabled**.
Logging	Specify the logging operation of this port. Notice that the logging message doesn't include the 4 bytes CRC. The allowed values are: **Enabled**: Frames received on the port are stored in the System Log. **Disabled**: Frames received on the port are not logged. The default value is **Disabled**. Note: The logging feature only works when the packet length is less than 1518 (without VLAN tags) and the System Log memory size and logging rate is limited.
Shutdown	Specify the port shut down operation of this port. The allowed values are: **Enabled**: If a frame is received on the port, the port will be disabled. **Disabled**: Port shut down is disabled. The default value is **Disabled**. *Note: The shutdown feature only works when the packet length is less than 1518 (without VLAN tags).*
State	Specify the port state of this port. The allowed values are: **Enabled**: To reopen ports by changing the volatile port configuration of the ACL user module.

	Disabled: To close ports by changing the volatile port configuration of the ACL user module. The default value is **Enabled**.
Counter	Counts the number of frames that match this ACE.

Rate Limiters

ACL rate limitation controls the traffic rate of packets that match Access Control List rules. It is used to limit or suppress excessive or unwanted traffic, protecting the device and network from congestion or denial-of-service conditions. Different rate limits can be configured for each *Rate Limiter ID* and then Assigned to an ACE.
ACL status can be viewed at Monitor→Security→Network→ACL Status.

ACL Rate Limiter Configuration

Rate Limiter ID	Rate	Unit
*	10	<>
1	10	pps
2	10	pps
3	10	pps
4	10	pps
5	10	pps
6	10	pps
7	10	pps
8	10	pps
9	10	pps
10	10	pps
11	10	pps
12	10	pps
13	10	pps
14	10	pps
15	10	pps
16	10	pps

Parameter	Description
Rate Limiter ID	The rate limiter ID for the settings contained in the same row and its range is **1** to **16**.
Rate	The valid rate is **0, 10, 20, 30, ..., 5000000** in pps or **0, 25, 50, 75, ..., 10000000** in kbps.
Unit	Specify the rate unit. The allowed values are: **pps**: packets per second. **kbps**: Kbits per second.

Access Control List

This page shows the Access Control List (ACL), which is made up of the ACEs defined on this switch. Each row describes the ACE that is defined. The maximum number of ACEs is 512 on each switch.
This table shows an overview of the currently configured ACEs. Click on the lowest plus sign to open the webpage to add a new ACE to the list (see below). The reserved ACEs used for internal protocol cannot be edited or deleted; the order sequence cannot be changed, and the priority is highest.
ACL status can be viewed at Monitor→Security→Network→ACL Status.

Access Control List Configuration

ACE	Ingress Port	Policy / Bitmask	Frame Type	Action	Rate Limiter	Port Redirect	Mirror	Counter	
1	All	Any	Any	Permit	Disabled	Disabled	Disabled	194	

Parameter	Description
ACE	Indicates the ACE ID.
Ingress Port	Indicates the ingress port of the ACE. Possible values are: **All**: The ACE will match all ingress port. **Port**: The ACE will match a specific ingress port.
Policy / Bitmask	Indicates the policy number and bitmask of the ACE.
Frame Type	Indicates the frame type of the ACE. Possible values are: **Any**: The ACE will match any frame type. **EType**: The ACE will match Ethernet Type frames. Note that an Ethernet Type based ACE will not get matched by IP and ARP frames. **ARP**: The ACE will match ARP/RARP frames. **IPv4**: The ACE will match all IPv4 frames. **IPv4/ICMP**: The ACE will match IPv4 frames with ICMP protocol. **IPv4/UDP**: The ACE will match IPv4 frames with UDP protocol. **IPv4/TCP**: The ACE will match IPv4 frames with TCP protocol. **IPv4/Other**: The ACE will match IPv4 frames, which are not ICMP/UDP/TCP. **IPv6**: The ACE will match all IPv6 standard frames.
Action	Indicates the forwarding action of the ACE. **Permit**: Frames matching the ACE may be forwarded and learned. **Deny**: Frames matching the ACE are dropped. **Filter**: Frames matching the ACE are filtered. The Filter action in ACL discards matching packets (drop), similar to Deny but without incrementing Deny statistics counters. It typically does not forward packets to the CPU, thereby maintaining hardware forwarding performance.
Rate Limiter	Indicates the rate limiter number of the ACE as defined in Configuration→Security→Network→ACL→ Rate Limiters. The allowed range is **1** to **16**. When **Disabled** is displayed, the rate limiter operation is disabled.
Port Redirect	Indicates the port redirect operation of the ACE. Frames matching the ACE are redirected to the port number. The allowed values are **Disabled** or a specific **port number**. When **Disabled** is displayed, the port redirect operation is disabled.
Mirror	Specify the mirror operation of this port. Frames matching the ACE are mirrored to the destination mirror port. The allowed values are: **Enabled**: Frames received on the port are mirrored. **Disabled**: Frames received on the port are not mirrored. The default value is "Disabled".
Counter	The counter indicates the number of times the ACE was hit by a frame.
Modification Buttons	You can modify each ACE (Access Control Entry) in the table using the following buttons: : Inserts a new ACE before the current row. : Edits the ACE row. : Moves the ACE up the list. : Moves the ACE down the list. : Deletes the ACE. : The lowest plus sign adds a new entry at the bottom of the ACE listings.

ACE Configuration

An ACE consists of several parameters. These parameters vary according to the frame type that you select. First select the ingress port for the ACE, and then select the frame type. Different parameter options are displayed depending on the frame type selected, and additional options become available when selecting *Specific* filters. A frame that hits this ACE matches the configuration that is defined here.

ACE Configuration

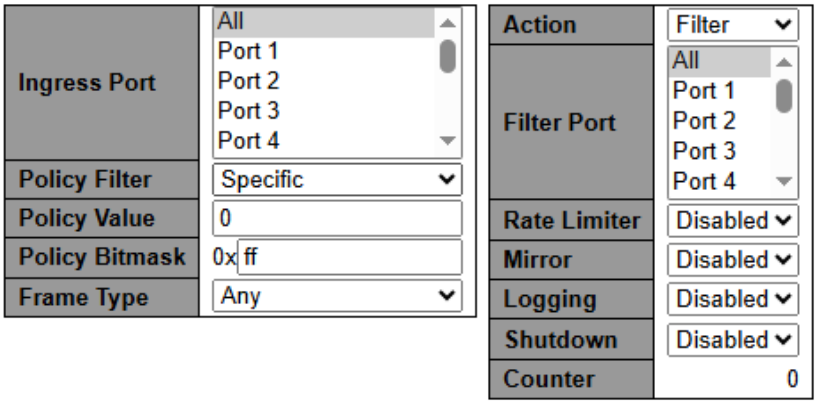

Parameter	Description
Ingress Port	Select the ingress port for which this ACE applies. **`All`**: The ACE applies to all port. **`Port n`**: The ACE applies to this port number, where *n* is the number of the switch port.
Policy Filter	Specify the policy number filter for this ACE. **`Any`**: No policy filter is specified. (policy filter status is "don't-care".) **`Specific`**: If you want to filter a specific policy with this ACE, choose this value. Two fields for entering an policy value and bitmask appear.
Policy Value	When "**`Specific`**" is selected for the policy filter, you can enter a specific policy value. The allowed range is **`0`** to **`255`**.
Policy Bitmask	When "**`Specific`**" is selected for the *policy filter*, you can enter a specific policy bitmask. The allowed range is **`0x0`** to **`0xff`**. Notice the usage of bitmask, if the binary bit value is "0", it means this bit is "don't-care". The real matched pattern is **`[policy_value & policy_bitmask]`**. For example, if the policy value is 3 and the policy bitmask is 0x10(bit 0 is "don't-care" bit), then policy 2 and 3 are applied to this rule.
Frame Type	Select the frame type for this ACE. These frame types are mutually exclusive. **`Any`**: Any frame can match this ACE. **`Ethernet Type`**: Only Ethernet Type frames can match this ACE. The IEEE 802.3 describes the value of Length/Type Field specifications to be greater than or equal to 1536 decimal (equal to 0600 hexadecimal) and the value should not be equal to 0x800(IPv4), 0x806(ARP) or 0x86DD(IPv6). **`ARP`**: Only ARP frames can match this ACE. Notice the ARP frames won't match the ACE with Ethernet type. **`IPv4`**: Only IPv4 frames can match this ACE. Notice the IPv4 frames won't match the ACE with Ethernet type. **`IPv6`**: Only IPv6 frames can match this ACE. Notice the IPv6 frames won't match the ACE with Ethernet type.
Action	Specify the action to take with a frame that hits this ACE. **`Permit`**: The frame that hits this ACE is granted permission for the ACE operation. **`Deny`**: The frame that hits this ACE is dropped. **`Filter`**: Frames matching the ACE are filtered.
Filter Port	When **`Filter`** is selected as an *Action*, you can select a filter port here.
Rate Limiter	Specify the rate limiter in number of base units. The allowed range is **`1`** to **`16`**. **`Disabled`** indicates that the rate limiter operation is disabled.
Port Redirect	Frames that hit the ACE are redirected to the port number specified here. The rate limiter will affect these ports. The allowed range is the same as the switch port number range. **`Disabled`** indicates that the port redirect operation is disabled and the specific port number of *'Port Redirect'* can't be set when action is permitted.
Mirror	Specify the mirror operation of this port. Frames matching the ACE are mirrored to the destination mirror port. The rate limiter will not affect frames on the mirror port. The allowed values are: **`Enabled`**: Frames received on the port are mirrored. **`Disabled`**: Frames received on the port are not mirrored. The default value is "**`Disabled`**".
Logging	Specify the logging operation of the ACE. Notice that the logging message doesn't include the 4 bytes CRC information. The allowed values are:

	Enabled: Frames matching the ACE are stored in the System Log. **Disabled**: Frames matching the ACE are not logged. Note: The logging feature only works when the packet length is less than 1518 (without VLAN tags) and the System Log memory size and logging rate is limited.
Shutdown	Specify the port shut down operation of the ACE. The allowed values are: **Enabled**: If a frame matches the ACE, the ingress port will be disabled. **Disabled**: Port shut down is disabled for the ACE. Note: The shutdown feature only works when the packet length is less than 1518 (without VLAN tags).
Counter	The counter indicates the number of times the ACE was hit by a frame.

VLAN Parameters

802.1Q Tagged	Any
VLAN ID Filter	Specific
VLAN ID	1
Tag Priority	Any

Parameter	Description
802.1Q Tagged	Specify whether frames can hit the action according to the 802.1Q tagged. The allowed values are: **Any**: Any value is allowed ("don't-care"). **Enabled**: Tagged frame only. **Disabled**: Untagged frame only. The default value is **Any**.
VLAN ID Filter	Specify the VLAN ID filter for this ACE. **Any**: No VLAN ID filter is specified. (VLAN ID filter status is "don't-care".) **Specific**: If you want to filter a specific VLAN ID with this ACE, choose this value. A field for entering a VLAN ID number appears.
VLAN ID	When "**Specific**" is selected for the *VLAN ID filter*, you can enter a specific VLAN ID number. The allowed range is **1** to **4095**. A frame that hits this ACE matches this VLAN ID value.
Tag Priority	Specify the tag priority for this ACE. A frame that hits this ACE matches this tag priority. The allowed number range is **0** to **7** or range **0-1**, **2-3**, **4-5**, **6-7**, **0-3** and **4-7**. The value **Any** means that no tag priority is specified (tag priority is "don't-care".)

MAC Parameters

SMAC Filter	Specific
SMAC Value	00-00-00-00-00-01
DMAC Filter	Specific
DMAC Value	00-00-00-00-00-02

Parameter	Description
SMAC Filter	Specify the source MAC filter for this ACE. **Any**: No SMAC filter is specified. (SMAC filter status is "don't-care".) **Specific**: If you want to filter a specific source MAC address with this ACE, choose this value. A field for entering a SMAC value appears. *(Only displayed when the frame type is Ethernet Type or ARP.)*
SMAC Value	When "**Specific**" is selected for the *SMAC filter*, you can enter a specific source MAC address. The legal format is "**-xx-xx-xx-xx-xx**" or "**xx.xx.xx.xx.xx.xx**" or "**xxxxxxxxxxxx**" " (x is a hexadecimal digit). A frame that hits this ACE matches this SMAC value.
DMAC Filter	Specify the destination MAC filter for this ACE. **Any**: No DMAC filter is specified. (DMAC filter status is "don't-care".) **MC**: Frame must be multicast. **BC**: Frame must be broadcast. **UC**: Frame must be unicast. **Specific**: If you want to filter a specific destination MAC address with this ACE, choose this value. A field for entering a DMAC value appears.

DMAC Value	When "**Specific**" is selected for the *DMAC filter*, you can enter a specific destination MAC address. The legal format is "**xx-xx-xx-xx-xx-xx**" or "**xx.xx.xx.xx.xx.xx**" or "**xxxxxxxxxxxx**" (x is a hexadecimal digit). A frame that hits this ACE matches this DMAC value.

Ethernet Type Parameters

EtherType Filter	Specific
Ethernet Type Value	0x FFFF

Parameter	Description
EtherType Filter	Specify the Ethernet type filter for this ACE. **Any**: No EtherType filter is specified (EtherType filter status is "don't-care"). **Specific**: If you want to filter a specific EtherType filter with this ACE, you can enter a specific EtherType value. A field for entering a EtherType value appears.
Ethernet Type Value	When "**Specific**" is selected for the *EtherType filter*, you can enter a specific EtherType value. The allowed range is **0x600** to **0xFFFF** but excluding 0x800(IPv4), 0x806(ARP) and 0x86DD(IPv6). A frame that hits this ACE matches this EtherType value.

ARP Parameters

ARP/RARP	Any	ARP Sender MAC Match	Any
Request/Reply	Any	RARP Target MAC Match	Any
Sender IP Filter	Network	IP/Ethernet Length	Any
Sender IP Address	0.0.0.0	IP	Any
Sender IP Mask	255.255.255.0	Ethernet	Any
Target IP Filter	Network		
Target IP Address	0.0.0.0		
Target IP Mask	255.255.255.0		

Parameter	Description
ARP/RARP	Specify the available ARP/RARP opcode (OP) flag for this ACE. **Any**: No ARP/RARP OP flag is specified. (OP is "don't-care".) **ARP**: Frame must have ARP opcode set to ARP. **RARP**: Frame must have RARP opcode set to RARP. **Other**: Frame has unknown ARP/RARP Opcode flag.
Request/Reply	Specify the available Request/Reply opcode (OP) flag for this ACE. **Any**: No Request/Reply OP flag is specified. (OP is "don't-care".) **Request**: Frame must have ARP Request or RARP Request OP flag set. **Reply**: Frame must have ARP Reply or RARP Reply OP flag.
Sender IP Filter	Specify the sender IP filter for this ACE. **Any**: No sender IP filter is specified. (Sender IP filter is "don't-care".) **Host**: Sender IP filter is set to Host. Specify the sender IP address in the SIP Address field that appears. **Network**: Sender IP filter is set to Network. Specify the sender IP address and sender IP mask in the SIP Address and SIP Mask fields that appear.
Sender IP Address	When **Host** or **Network** is selected for the *sender IP filter*, you can enter a specific sender IP address in dotted decimal notation. Notice the invalid IP address configuration is acceptable too, for example, **0.0.0.0**. Normally, an ACE with invalid IP address will explicitly adding deny action.
Sender IP Mask	When "**Network**" is selected for the sender IP filter, you can enter a specific sender IP mask in dotted decimal notation.
Target IP Filter	Specify the target IP filter for this specific ACE. **Any**: No target IP filter is specified. (Target IP filter is "don't-care".) **Host**: Target IP filter is set to Host. Specify the target IP address in the *Target IP Address field*. **Network**: Target IP filter is set to Network. Specify the target IP address and target IP mask in the Target IP Address and Target IP Mask fields.
Target IP Address	When **Host** or **Network** is selected for the *target IP filter*, you can enter a specific target IP address in dotted decimal notation. Notice the invalid IP address configuration is acceptable too, for example, **0.0.0.0**. Normally, an ACE with invalid IP address will explicitly adding deny action.

Target IP Mask	When **Network** is selected for the *target IP filter*, you can enter a specific target IP mask in dotted decimal notation.
ARP Sender MAC Match	Specify whether frames can hit the action according to their sender hardware address field (SHA) settings. **0**: ARP frames where SHA is not equal to the SMAC address. **1**: ARP frames where SHA is equal to the SMAC address. **Any**: Any value is allowed ("don't-care").
RARP Target MAC Match	Specify whether frames can hit the action according to their target hardware address field (THA) settings. **0**: RARP frames where THA is not equal to the target MAC address. **1**: RARP frames where THA is equal to the target MAC address. **Any**: Any value is allowed ("don't-care").
IP/Ethernet Length	Specify whether frames can hit the action according to their ARP/RARP hardware address length (HLN) and protocol address length (PLN) settings. **0**: ARP/RARP frames where the HLN is not equal to Ethernet (0x06) or the (PLN) is not equal to IPv4 (0x04). **1**: ARP/RARP frames where the HLN is equal to Ethernet (0x06) and the (PLN) is equal to IPv4 (0x04). **Any**: Any value is allowed ("don't-care").
IP	Specify whether frames can hit the action according to their ARP/RARP hardware address space (HRD) settings. **0**: ARP/RARP frames where the HLD is not equal to Ethernet (1). **1**: ARP/RARP frames where the HLD is equal to Ethernet (1). **Any**: Any value is allowed ("don't-care").
Ethernet	Specify whether frames can hit the action according to their ARP/RARP protocol address space (PRO) settings. **0**: ARP/RARP frames where the PRO is not equal to IP (0x800). **1**: ARP/RARP frames where the PRO is equal to IP (0x800). **Any**: Any value is allowed ("don't-care").

IP Parameters

IP Protocol Filter	Other
IP Protocol Value	255
IP TTL	Any
IP Fragment	Any
IP Option	Any
SIP Filter	Network
SIP Address	0.0.0.0
SIP Mask	255.255.255.0
DIP Filter	Network
DIP Address	0.0.0.0
DIP Mask	255.255.255.0

Parameter	Description
IP Protocol Filter	Specify the IP protocol filter for this ACE. **Any**: No IP protocol filter is specified ("don't-care"). **Specific**: If you want to filter a specific IP protocol filter with this ACE, choose this value. A field for entering an IP protocol filter appears. **ICMP**: Select ICMP to filter IPv4 ICMP protocol frames. Extra fields for defining ICMP parameters will appear. These fields are explained later in this help file. **UDP**: Select UDP to filter IPv4 UDP protocol frames. Extra fields for defining UDP parameters will appear. These fields are explained later in this help file. **TCP**: Select TCP to filter IPv4 TCP protocol frames. Extra fields for defining TCP parameters will appear. These fields are explained later in this help file.
IP Protocol Value	When "**Specific**" is selected for the *IP protocol filter*, you can enter a specific value. The allowed range is **0** to **255**. A frame that hits this ACE matches this IP protocol value.
IP TTL	Specify the Time-to-Live settings for this ACE. **zero**: IPv4 frames with a Time-to-Live field greater than zero must not be able to match this entry.

	non-zero: IPv4 frames with a Time-to-Live field greater than zero must be able to match this entry. **Any**: Any value is allowed ("don't-care").
IP Fragment	Specify the fragment offset settings for this ACE. This involves the settings for the More Fragments (MF) bit and the Fragment Offset (FRAG OFFSET) field for an IPv4 frame. **No**: IPv4 frames where the MF bit is set or the FRAG OFFSET field is greater than zero must not be able to match this entry. **Yes**: IPv4 frames where the MF bit is set or the FRAG OFFSET field is greater than zero must be able to match this entry. **Any**: Any value is allowed ("don't-care").
IP Option	Specify the options flag setting for this ACE. **No**: IPv4 frames where the options flag is set must not be able to match this entry. **Yes**: IPv4 frames where the options flag is set must be able to match this entry. **Any**: Any value is allowed ("don't-care").
SIP Filter	Specify the source IP filter for this ACE. **Any**: No source IP filter is specified. (Source IP filter is "don't-care".) **Host**: Source IP filter is set to Host. Specify the source IP address in the SIP Address field that appears. **Network**: Source IP filter is set to Network. Specify the source IP address and source IP mask in the SIP Address and SIP Mask fields that appear.
SIP Address	When "**Host**" or "**Network**" is selected for the *source IP filter*, you can enter a specific SIP address in dotted decimal notation. Notice the invalid IP address configuration is acceptable too, for example, **0.0.0.0**. Normally, an ACE with invalid IP address will explicitly adding deny action.
SIP Mask	When "**Network**" is selected for the *source IP filter*, you can enter a specific SIP mask in dotted decimal notation.
DIP Filter	Specify the destination IP filter for this ACE. **Any**: No destination IP filter is specified. (Destination IP filter is "don't-care".) **Host**: Destination IP filter is set to Host. Specify the destination IP address in the DIP Address field that appears. **Network**: Destination IP filter is set to Network. Specify the destination IP address and destination IP mask in the DIP Address and DIP Mask fields that appear.
DIP Address	When "**Host**" or "**Network**" is selected for the *destination IP filter*, you can enter a specific DIP address in dotted decimal notation. Notice the invalid IP address configuration is acceptable too, for example, **0.0.0.0**. Normally, an ACE with invalid IP address will explicitly adding deny action.
DIP Mask	When "**Network**" is selected for the *destination IP filter*, you can enter a specific DIP mask in dotted decimal notation.

IPv6 Parameters

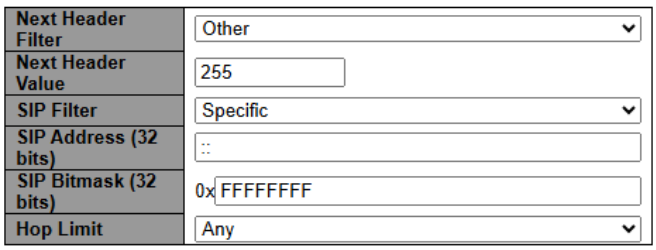

Parameter	Description
Next Header Filter	Specify the IPv6 next header filter for this ACE. **Any**: No IPv6 next header filter is specified ("don't-care"). **Specific**: If you want to filter a specific IPv6 next header filter with this ACE, choose this value. A field for entering an IPv6 next header filter appears. **ICMP**: Select ICMP to filter IPv6 ICMP protocol frames. Extra fields for defining ICMP parameters will appear. These fields are explained later in this help file. **UDP**: Select UDP to filter IPv6 UDP protocol frames. Extra fields for defining UDP parameters will appear. These fields are explained later in this help file. **TCP**: Select TCP to filter IPv6 TCP protocol frames. Extra fields for defining TCP parameters will appear. These fields are explained later in this help file.
Next Header Value	When "**Specific**" is selected for the *IPv6 next header filter*, you can enter a specific value. The allowed range is **0** to **255**. A frame that hits this ACE matches this IPv6 protocol value.

SIP Filter	Specify the source IPv6 filter for this ACE. **Any**: No source IPv6 filter is specified. (Source IPv6 filter is "don't-care".) **Specific**: Source IPv6 filter is set to Network. Specify the source IPv6 address and source IPv6 mask in the SIP Address fields that appear.
SIP Address	When "**Specific**" is selected for the *source IPv6 filter*, you can enter a specific SIPv6 address. The field only supported last 32 bits for IPv6 address.
SIP BitMask	When "**Specific**" is selected for the *source IPv6 filter*, you can enter a specific SIPv6 mask. The field only supported last 32 bits for IPv6 address. Notice the usage of bitmask, if the binary bit value is "0", it means this bit is "don't-care". The real matched pattern is **[sipv6_address & sipv6_bitmask]** (last 32 bits). For example, if the SIPv6 address is **2001::3** and the SIPv6 bitmask is **0xFFFFFFFE** (bit 0 is "don't-care" bit), then SIPv6 address **2001::2** and **2001::3** are applied to this rule.
Hop Limit	Specify the hop limit settings for this ACE. **zero**: IPv6 frames with a hop limit field greater than zero must not be able to match this entry. **non-zero**: IPv6 frames with a hop limit field greater than zero must be able to match this entry. **Any**: Any value is allowed ("don't-care").

ICMP Parameters

ICMP Type Filter	Specific
ICMP Type Value	255
ICMP Code Filter	Specific
ICMP Code Value	255

Parameter	Description
ICMP Type Filter	Specify the ICMP filter for this ACE. **Any**: No ICMP filter is specified (ICMP filter status is "don't-care"). **Specific**: If you want to filter a specific ICMP filter with this ACE, you can enter a specific ICMP value. A field for entering an ICMP value appears.
ICMP Type Value	When "**Specific**" is selected for the *ICMP type filter*, you can enter a specific value. The allowed range is **0** to **255**. A frame that hits this ACE matches this ICMP value.
ICMP Code Filter	Specify the ICMP code filter for this ACE. **Any**: No ICMP code filter is specified (ICMP code filter status is "don't-care"). **Specific**: If you want to filter a specific ICMP code filter with this ACE, you can enter a specific ICMP code value. A field for entering an ICMP code value appears.
ICMP Code Value	When "**Specific**" is selected for the *ICMP code filter*, you can enter a specific ICMP code value. The allowed range is **0** to **255**. A frame that hits this ACE matches this ICMP code value.

TCP Parameters

Source Port Filter	Specific
Source Port No.	0
Dest. Port Filter	Specific
Dest. Port No.	0
TCP FIN	Any
TCP SYN	Any
TCP RST	Any
TCP PSH	Any
TCP ACK	Any
TCP URG	Any

UDP Parameters

Source Port Filter	Specific
Source Port No.	0
Dest. Port Filter	Specific
Dest. Port No.	0

Parameter	Description
TCP/UDP Source Filter	Specify the TCP/UDP source filter for this ACE. **Any**: No TCP/UDP source filter is specified (TCP/UDP source filter status is "don't-care"). **Specific**: If you want to filter a specific TCP/UDP source filter with this ACE, you can enter a specific TCP/UDP source value. A field for entering a TCP/UDP source value appears. **Range**: If you want to filter a specific TCP/UDP source range filter with this ACE, you can enter a specific TCP/UDP source range value. A field for entering a TCP/UDP source value appears.

TCP/UDP Source No.	When "**Specific**" is selected for the *TCP/UDP source filter*, you can enter a specific TCP/UDP source value. The allowed range is **0** to **65535**. A frame that hits this ACE matches this TCP/UDP source value.
TCP/UDP Source Range	When "**Range**" is selected for the *TCP/UDP source filter*, you can enter a specific TCP/UDP source range value. The allowed range is **0** to **65535**. A frame that hits this ACE matches this TCP/UDP source value.
TCP/UDP Destination Filter	Specify the TCP/UDP destination filter for this ACE. **Any**: No TCP/UDP destination filter is specified (TCP/UDP destination filter status is "don't-care"). **Specific**: If you want to filter a specific TCP/UDP destination filter with this ACE, you can enter a specific TCP/UDP destination value. A field for entering a TCP/UDP destination value appears. **Range**: If you want to filter a specific range TCP/UDP destination filter with this ACE, you can enter a specific TCP/UDP destination range value. A field for entering a TCP/UDP destination value appears.
TCP/UDP Destination Number	When "**Specific**" is selected for the *TCP/UDP destination filter*, you can enter a specific TCP/UDP destination value. The allowed range is **0** to **65535**. A frame that hits this ACE matches this TCP/UDP destination value.
TCP/UDP Destination Range	When "**Range**" is selected for the *TCP/UDP destination filter*, you can enter a specific TCP/UDP destination range value. The allowed range is **0** to **65535**. A frame that hits this ACE matches this TCP/UDP destination value.
TCP FIN	Specify the TCP "No more data from sender" (FIN) value for this ACE. **0**: TCP frames where the FIN field is set must not be able to match this entry. **1**: TCP frames where the FIN field is set must be able to match this entry. **Any**: Any value is allowed ("don't-care").
TCP SYN	Specify the TCP "Synchronize sequence numbers" (SYN) value for this ACE. **0**: TCP frames where the SYN field is set must not be able to match this entry. **1**: TCP frames where the SYN field is set must be able to match this entry. **Any**: Any value is allowed ("don't-care").
TCP RST	Specify the TCP "Reset the connection" (RST) value for this ACE. **0**: TCP frames where the RST field is set must not be able to match this entry. **1**: TCP frames where the RST field is set must be able to match this entry. **Any**: Any value is allowed ("don't-care").
TCP PSH	Specify the TCP "Push Function" (PSH) value for this ACE. **0**: TCP frames where the PSH field is set must not be able to match this entry. **1**: TCP frames where the PSH field is set must be able to match this entry. **Any**: Any value is allowed ("don't-care").
TCP ACK	Specify the TCP "Acknowledgment field significant" (ACK) value for this ACE. **0**: TCP frames where the ACK field is set must not be able to match this entry. **1**: TCP frames where the ACK field is set must be able to match this entry. **Any**: Any value is allowed ("don't-care").
TCP URG	Specify the TCP "Urgent Pointer field significant" (URG) value for this ACE. **0**: TCP frames where the URG field is set must not be able to match this entry. **1**: TCP frames where the URG field is set must be able to match this entry. **Any**: Any value is allowed ("don't-care").

4.12.2.4 IP Source Guard

IP Source Guard is a secure feature used to restrict IP traffic on DHCP snooping untrusted ports by filtering traffic based on the DHCP Snooping Table or manually configured IP Source Bindings. It helps prevent IP spoofing attacks when a host tries to spoof and use the IP address of another host.

IP Source Guard can be viewed at Monitor→Security→Network→IP Source Guard.

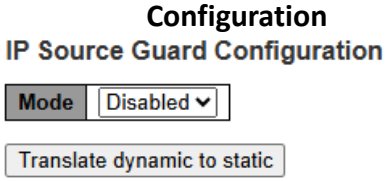

Port Mode Configuration

Port	Mode	Max Dynamic Clients
*	<>	<>
1	Disabled	Unlimited
2	Disabled	Unlimited
3	Disabled	Unlimited
4	Disabled	Unlimited
5	Disabled	Unlimited
6	Disabled	Unlimited
7	Disabled	Unlimited
8	Disabled	Unlimited
9	Disabled	Unlimited
10	Disabled	Unlimited
11	Disabled	Unlimited
12	Disabled	Unlimited

Parameter	Description
Mode of IP Source Guard Configuration	Enable the Global IP Source Guard or disable the Global IP Source Guard. All configured ACEs will be lost when the mode is enabled.
Port Mode Configuration	Specify IP Source Guard is enabled on which ports. Only when both Global Mode and Port Mode on a given port are enabled, IP Source Guard is enabled on this given port.
Max Dynamic Clients	Specify the maximum number of dynamic clients that can be learned on given port. This value can be 0, 1, 2 or unlimited. If the port mode is enabled and the value of max dynamic client is equal to 0, it means only allow the IP packets forwarding that are matched in static entries on the specific port.

Static Table

This page shows the static IP Source Guard rules. The static table contains pre-configured IP-to-MAC bindings that are enforced on trusted ports. Traffic is permitted only if the source IP and MAC address match an entry in the table, preventing IP spoofing and ensuring consistent device identity.
The maximum number of rules is **112** on the switch.
IP Source Guard can be viewed at Monitor→Security→Network→IP Source Guard.

Static IP Source Guard Table

Delete	Port	VLAN ID	IP Address	MAC address
Delete	1			

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
Port Mode Configuration	The logical port for the settings.
VLAN ID	The VLAN ID for the settings.
IP Address	Allowed Source IP address.
MAC address	Allowed Source MAC address.

4.12.2.5 IPv6 Source Guard

The IPv6 Source Guard uses identical parameters to the IPv4 Source Guard. For a detailed description, please refer to Configuration→Security→Network→IP Source Guard.

4.12.2.6 ARP Inspection

ARP Inspection is a secure feature. Several types of attacks can be launched against a host or devices connected to Layer 2 networks by "poisoning" the ARP caches. This feature is used to block such attacks. Only valid ARP requests and responses can go through the switch device.
ARP Inspection can be viewed at Monitor→Security→Network→ARP Inspection.

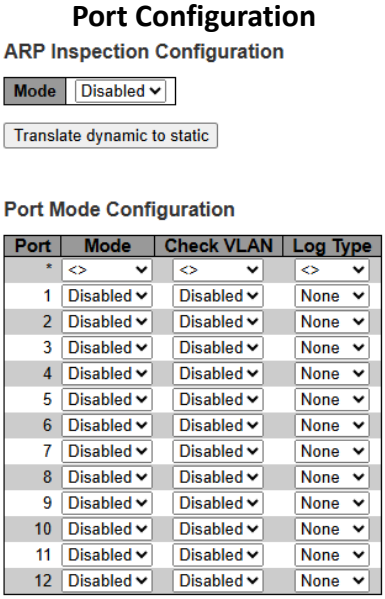

Parameter	Description
Global Mode	Enable the Global ARP Inspection or disable the Global ARP Inspection.
Port	Specifies the individual port settings.
Port Mode	Specify ARP Inspection is enabled on which ports. Only when both Global Mode and Port Mode on a given port are enabled, ARP Inspection is enabled on this given port. Possible modes are: **Enabled**: Enable ARP Inspection operation. **Disabled**: Disable ARP Inspection operation.
Check VLAN	If you want to inspect the VLAN configuration, you have to enable the setting of "Check VLAN". The default setting of "Check VLAN" is disabled. When the setting of "Check VLAN" is disabled, the log type of ARP Inspection will refer to the port setting. And the setting of "Check VLAN" is enabled, the log type of ARP Inspection will refer to the VLAN setting. Possible setting of "Check VLAN" are: **Enabled**: Enable check VLAN operation. **Disabled**: Disable check VLAN operation.
IP Address	Only the Global Mode and Port Mode on a given port are enabled, and the setting of "Check VLAN" is disabled, the log type of ARP Inspection will refer to the port setting. There are four log types and possible types are: **None**: Log nothing. **Deny**: Log denied entries. **Permit**: Log permitted entries. **ALL**: Log all entries.

VLAN Configuration

Specify which VLANs ARP Inspection is enabled. First, you must enable the port setting at Configuration→Security→Network→ARP Inspection→Port Configuration. Only when both Global Mode and Port Mode on a given port are enabled, ARP Inspection is enabled on this given port.

Each page shows up to 9999 entries from the VLAN table; default being 20, selected through the "entries per page" input field. When first visited, the web page will show the first 20 entries from the beginning of the VLAN Table. The first displayed will be the one with the lowest VLAN ID found in the VLAN Table.

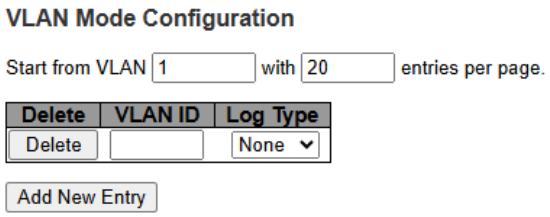

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
VLAN ID	Specify which VLAN will be inspected.
Log Type	Specify which events will be logged. Possible types are: **None**: Log nothing. **Deny**: Log denied entries. **Permit**: Log permitted entries. **ALL**: Log all entries.

Static Table

Configuration of the static ARP Inspection rules. The maximum number of rules is **256** on the switch.

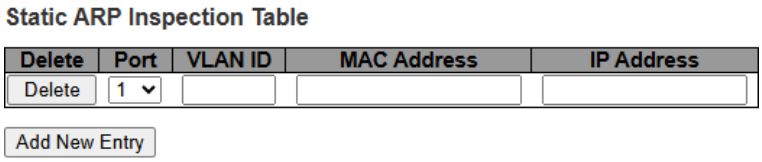

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
Port	The logical port for the settings.
VLAN ID	Specify which VLAN will be inspected.
MAC Address	Allowed Source MAC address in ARP request packets.
IP Address	Allowed Source IP address in ARP request packets.

Dynamic Table

Entries in the Dynamic ARP Inspection Table are shown on this page. The Dynamic ARP Inspection Table contains up to **256** entries, and is sorted first by port, then by VLAN ID, then by MAC address, and then by IP address. All dynamic entries are learned from DHCP Snooping.

Each page shows up to 99 entries from the Dynamic ARP Inspection table, default being 20, selected through the "entries per page" input field. When first visited, the web page will show the first 20 entries from the beginning of the Dynamic ARP Inspection Table.
The "Start from port address", "VLAN", "MAC address" and "IP address" input fields allow the user to select the starting point in the Dynamic ARP Inspection Table.

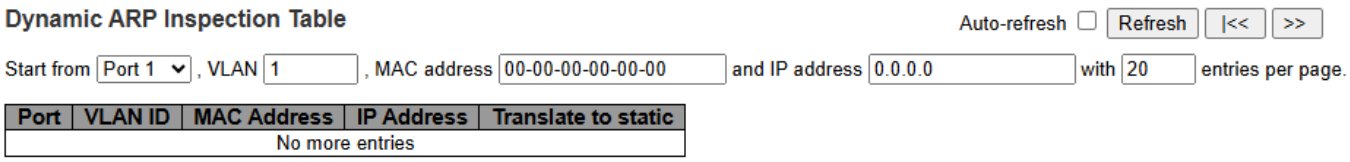

Parameter	Description
Port	Switch Port Number for which the entries are displayed.
VLAN ID	VLAN-ID in which the ARP traffic is permitted.
MAC Address	User MAC address of the entry.
IP Address	User IP address of the entry.
Translate to static	Select the checkbox to translate the entry to static entry.

4.12.3 AAA

This section describes the configuration of remote servers for authentication, authorization, and accounting on the Ethernet switch.

4.12.3.1 RADIUS

RADIUS is an acronym for **R**emote **A**uthentication **D**ial **I**n **U**ser **S**ervice. It is a networking protocol that provides centralized access, authorization, and accounting management for people or computers to connect and use a network service. Up to 5 RADIUS servers can be configured.
RADIUS can be viewed at Monitor→Security→AAA.

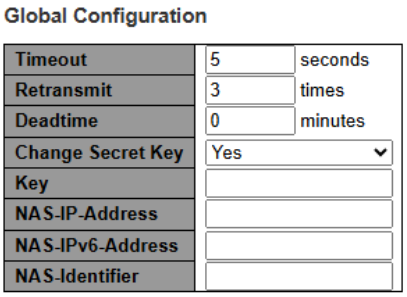

Parameter	Description
Timeout	Timeout is the number of seconds, in the range **1 to 1000**, to wait for a reply from a RADIUS server before retransmitting the request.
Retransmit	Retransmit is the number of times, in the range **1 to 1000**, a RADIUS request is retransmitted to a server that is not responding. If the server has not responded after the last retransmit it is considered to be dead.
Deadtime	Deadtime, which can be set to a number between **0 to 1440** minutes, is the period during which the switch will not send new requests to a server that has failed to respond to a previous request. This will stop the switch from continually trying to contact a server that it has already determined as dead.

	Setting the Deadtime to a value greater than 0 (zero) will enable this feature, but only if more than one server has been configured.
Change Secret Key	Specify to change the secret key or not. When "Yes" is selected for the option, you can change the secret key - up to 63 characters long - shared between the RADIUS server and the switch.
NAS-IP-Address (Attribute 4)	The IPv4 address to be used as attribute 4 in RADIUS Access-Request packets. If this field is left blank, the IP address of the outgoing interface is used.
NAS-IPv6-Address (Attribute 95)	The IPv6 address to be used as attribute 95 in RADIUS Access-Request packets. If this field is left blank, the IP address of the outgoing interface is used.
NAS-Identifier (Attribute 32)	The identifier - up to 253 characters long - to be used as attribute 32 in RADIUS Access-Request packets. If this field is left blank, the NAS-Identifier is not included in the packet.

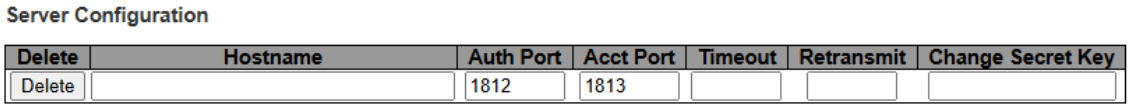

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
Hostname	The IPv4/IPv6 address or hostname of the RADIUS server.
Auth Port	The UDP port to use on the RADIUS server for authentication. Set to 0 to disable authentication.
Acct Port	The UDP port to use on the RADIUS server for accounting. Set to 0 to disable accounting.
Timeout	This optional setting overrides the global timeout value. Leaving it blank will use the global timeout value.
Retransmit	This optional setting overrides the global retransmit value. Leaving it blank will use the global retransmit value.
Change Secret Key	Specify to change the secret key or not. When the checkbox is checked, you can change the setting overrides the global key. Leaving it blank will use the global key.

4.12.3.2 TACACS+

TACACS+ is an acronym for **T**erminal **A**ccess **C**ontroller **A**ccess **C**ontrol **S**ystem **P**lus. It is a networking protocol which provides access control for routers, network access servers and other networked computing devices via one or more centralized servers. TACACS+ provides separate authentication, authorization and accounting services. Up to 5 TACACS+ servers can be configured.

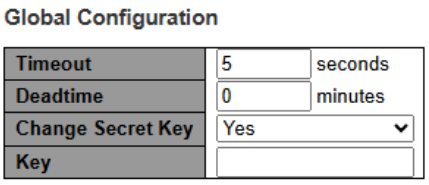

Parameter	Description
Timeout	Timeout is the number of seconds, in the range **1 to 1000**, to wait for a reply from a TACACS+ server before it is considered to be dead.
Deadtime	Deadtime, which can be set to a number between **0 to 1440** minutes, is the period during which the switch will not send new requests to a server that has failed to respond to a previous request. This will stop the switch from continually trying to contact a server that it has already determined as dead. Setting the Deadtime to a value greater than 0 (zero) will enable this feature, but only if more than one server has been configured.
Change Secret Key	Specifies to change the secret key or not. When "Yes" is selected for the option, you can change the secret key - up to 63 characters long - shared between the TACACS+ server and the switch.

	When the "No" option is selected and no global key is currently configured, the "Key" input field is blank. If a key is configured, the "Key" input field contains asterisks.

Server Configuration

Delete	Hostname	Port	Timeout	Change Secret Key
Delete		49		

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
Hostname	The IPv4/IPv6 address or hostname of the TACACS+ server.
Port	The TCP port to use on the TACACS+ server for authentication.
Timeout	This optional setting overrides the global timeout value. Leaving it blank will use the global timeout value.
Change Secret Key	Specify to change the secret key or not. When the checkbox is checked, you can change the setting overrides the global key. Leaving it blank will use the global key. When the checkbox is unchecked, the field contains asterisks if a key is configured, otherwise it is blank.

4.13 Aggregation

Port or Link Aggregation uses multiple ports in parallel to increase the link speed beyond the limits of a port and to increase the redundancy for higher availability.

4.13.1 Common

This page is used to configure the Aggregation hash mode. This mode applies to the whole network element. View Aggregation status at Monitor→Aggregation→Status.

Common Aggregation Configuration

Hash Code Contributors	
Source MAC Address	☑
Destination MAC Address	☐
IP Address	☐
TCP/UDP Port Number	☐

Parameter	Description
Source MAC Address	The Source MAC address can be used to calculate the destination port for the frame. Check to enable the use of the Source MAC address, or uncheck to disable. By default, Source MAC Address is enabled.
Destination MAC Address	The Destination MAC Address can be used to calculate the destination port for the frame. Check to enable the use of the Destination MAC Address, or uncheck to disable. By default, Destination MAC Address is disabled.
IP Address	The IP address can be used to calculate the destination port for the frame. Check to enable the use of the IP Address, or uncheck to disable. By default, IP Address is enabled.
TCP/UDP Port Number	The TCP/UDP port number can be used to calculate the destination port for the frame. Check to enable the use of the TCP/UDP Port Number, or uncheck to disable. By default, TCP/UDP Port Number is enabled.

4.13.2 Groups

Aggregation groups combine multiple physical ports into one logical interface for increased bandwidth and redundancy. They define the member ports and operational parameters used for link aggregation, with or without LACP.
View Aggregation status at Monitor→Aggregation→Status.

Aggregation Group Configuration

	Port Members												Group Configuration		
Group ID	1	2	3	4	5	6	7	8	9	10	11	12	Mode	Revertive	Max Bundle
Normal	◉	◉	◉	◉	◉	◉	◉	◉	◉	◉	◉	◉			
1	○	○	○	○	○	○	○	○	○	○	○	○	Disabled	☑	12
2	○	○	○	○	○	○	○	○	○	○	○	○	Disabled	☑	12
3	○	○	○	○	○	○	○	○	○	○	○	○	Disabled	☑	12
4	○	○	○	○	○	○	○	○	○	○	○	○	Disabled	☑	12
5	○	○	○	○	○	○	○	○	○	○	○	○	Disabled	☑	12
6	○	○	○	○	○	○	○	○	○	○	○	○	Disabled	☑	12

Parameter	Description
Group ID	Indicates the aggregation group ID for the settings contained in the same row. Group ID "Normal" indicates there is no aggregation. Only one group ID is valid per port.
Port Members	Each switch port is listed for each group ID. Select a radio button to include a port in an aggregation, or clear the radio button to remove the port from the aggregation. By default, no ports belong to any aggregation group. Only full duplex ports can join an aggregation and ports must be in the same speed in each group.
Mode	This parameter determines the mode for the aggregation group. **Disabled**: The group is disabled. **Static**: The group operates in static aggregation mode. **LACP (Active)**: The group operates in LACP active aggregation mode. See IEEE 801.AX-2014, section 6.4.1 for details. **LACP (Passive)**: The group operates in LACP passive aggregation mode. See IEEE 801.AX-2014, section 6.4.1 for details.
Revertive	This parameter only applies to LACP-enabled groups. It determines if the group will perform automatic link (re-)calculation when links with higher priority becomes available.
Max Bundle	This parameter only applies to LACP-enabled groups. It determines the maximum number of active bundled LACP ports allowed in an aggregation.

4.13.3 LACP

The Link Aggregation Control Protocol (LACP) is an IEEE 802.3AX standard protocol which automatically manages the bundling of multiple physical links into a single logical link. It provides dynamic link aggregation, load balancing, and redundancy, ensuring efficient use of bandwidth and failover in case of link failure.
View LACP status and statistics at Monitor→Aggregation→LACP.

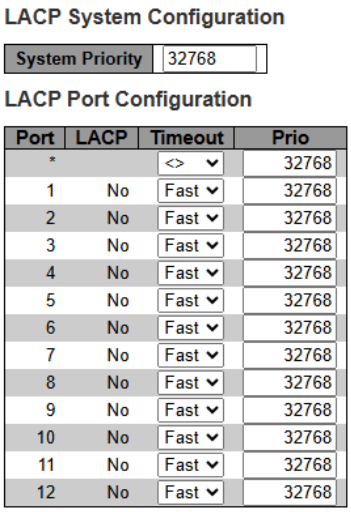

Parameter	Description
System Priority	Determines which switch acts as the master controller (Actor) during LACP negotiation between multiple switches. Lower values (**1-65535**, default 32768) have higher priority. The switch with the lowest System Priority + MAC address combination controls which ports become active in the LAG
Port	The switch port number.
LACP	Show whether LACP is currently enabled on this switch port.
Timeout	The *Timeout* controls the period between BPDU transmissions. **Fast** will transmit LACP packets each second, while **Slow** will wait for 30 seconds before sending a LACP packet.
Prio	The *Prio* controls the priority of the port, range **1-65535**. If the LACP partner wants to form a larger group than is supported by this device then this parameter will control which ports will be active and which ports will be in a backup role. Lower number means greater priority.
Max Bundle	This parameter only applies to LACP-enabled groups. It determines the maximum number of active bundled LACP ports allowed in an aggregation.

4.14 Link OAM

Link OAM (Operation Administration and Maintenance) is a protocol described in ITU-T Y.1731 used to implement carrier Ethernet functionality. Link OAM provides tools to monitor, troubleshoot, and maintain point-to-point Ethernet links. It supports fault detection, performance monitoring, and loopback testing, enabling rapid identification and isolation of link-level issues. MEP functionalities like CC and RDI are based on this.

4.14.1 Port Settings

View Link OAM port status at Monitor→Link OAM→Port Status.

Link OAM Port Configuration

Port	OAM Enabled	OAM Mode	Loopback Support	Link Monitor Support	MIB Retrieval Support	Loopback Operation
*	☐	<>	☐	☑	☐	☐
1	☐	Passive	☐	☑	☐	☐
2	☐	Passive	☐	☑	☐	☐
3	☐	Passive	☐	☑	☐	☐
4	☐	Passive	☐	☑	☐	☐
5	☐	Passive	☐	☑	☐	☐
6	☐	Passive	☐	☑	☐	☐
7	☐	Passive	☐	☑	☐	☐
8	☐	Passive	☐	☑	☐	☐
9	☐	Passive	☐	☑	☐	☐
10	☐	Passive	☐	☑	☐	☐
11	☐	Passive	☐	☑	☐	☐
12	☐	Passive	☐	☑	☐	☐

Parameter	Description
Port	The switch port number.
OAM Enabled	Controls whether Link OAM is enabled on this switch port. Enabling Link OAM provides the network operators the ability to monitor the health of the network and quickly determine the location of failing links or fault conditions.
OAM Mode	Configures the OAM Mode as Active or Passive. The default mode is Passive. **Active mode:** DTE's configured in Active mode initiate the exchange of Information OAMPDUs as defined by the Discovery process. Once the Discovery process completes, Active DTE's are permitted to send any OAMPDU while connected to a remote OAM peer entity in Active mode. Active DTE's operate in a limited respect if the remote OAM entity is operating in Passive mode. Active devices should not respond to OAM remote loopback commands and variable requests from a Passive peer. **Passive mode:** DTE's configured in Passive mode do not initiate the Discovery process. Passive DTE's react to the initiation of the Discovery process by the remote DTE. This eliminates the possibility of passive to passive links. Passive DTE's shall not send Variable Request or Loopback Control OAMPDUs.
Loopback Support	Controls whether the loopback support is enabled for the switch port. Link OAM remote loopback can be used for fault localization and link performance testing. Enabling the loopback support will allow the DTE to execute the remote loopback command that helps in the fault detection.
Link Monitor Support	Controls whether the Link Monitor support is enabled for the switch port. On enabling the Link Monitor support, the DTE supports event notification that permits the inclusion of diagnostic information.
MIB Retrieval Support	Controls whether the MIB Retrieval Support is enabled for the switch port. On enabling the MIB retrieval support, the DTE supports polling of various Link OAM based MIB variables' contents.
Loopback Operation	If the Loopback support is enabled, enabling this field will start a loopback operation for the port.

4.14.2 Event Settings

Link OAM event can be viewed at Monitor→Link OAM→Event Status.

Link Event Configuration for Port 1 Port 1

Event Name	Error Window	Error Threshold
Error Frame Event	1	1
Symbol Period Error Event	1	1
Seconds Summary Event	60	1

Parameter	Description

Port	The switch port number.
Event Name	Name of the Link Event which is being configured.
Error Window	Represents the window period in the order of 1 sec for the observation of various link events.
Error Threshold	Represents the threshold value for the window period for the appropriate Link event so as to notify the peer of this error.
Error Frame Event	The Errored Frame Event counts the number of errored frames detected during the specified period. The period is specified by a time interval (Window in order of 1 sec). This event is generated if the errored frame count is equal to or greater than the specified threshold for that period (Period Threshold). Errored frames are frames that had transmission errors as detected at the Media Access Control sublayer. Error Window for 'Error Frame Event' must be an integer value between **1-60** and its default value is '1'. Whereas Error Threshold must be between **0-4294967295** and its default value is '1'.
Symbol Period Error Event	The Errored Symbol Period Event counts the number of symbol errors that occurred during the specified period. The period is specified by the number of symbols that can be received in a time interval on the underlying physical layer. This event is generated if the symbol error count is equal to or greater than the specified threshold for that period. Error Window for 'Symbol Period Error Event' must be an integer value between **1-60** and its default value is '1'. Whereas Error Threshold must be between **0-4294967295** and its default value is '1'.
Seconds Summary Event	The Errored Frame Seconds Summary Event TLV counts the number of errored frame seconds that occurred during the specified period. The period is specified by a time interval. This event is generated if the number of errored frame seconds is equal to or greater than the specified threshold for that period. An errored frame second is a one second interval wherein at least one frame error was detected. Errored frames are frames that had transmission errors as detected at the Media Access Control sublayer. Error Window for 'Seconds Summary Event' must be an integer value between **10-900** and its default value is '60'. Whereas Error Threshold must be between **0-65535** and its default value is '1'.

4.15 Loop Protection

The Loop Protection feature prevents Layer 2 loops caused by incorrect cabling or wiring mistakes.
Loop Protection can be viewed at Monitor→Loop Protection.

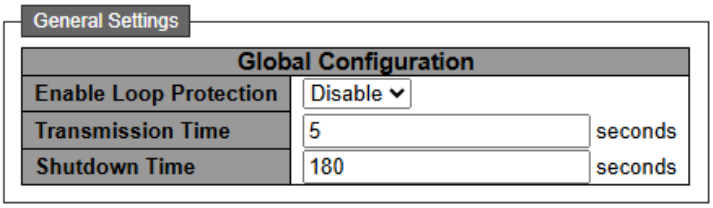

Parameter	Description
Enable Loop Protection	Enable Loop Protection
Transmission Time	The interval between each loop protection PDU sent on each port. Valid values are **1 to 10** seconds. Default value is 5 seconds.
Shutdown Time	The period (in seconds) for which a port will be kept disabled in the event of a loop is detected (and the port action shuts down the port). Valid values are **0 to 604800** seconds (7 days). A value of zero will keep a port disabled (until next device restart). Default value is 180 seconds.

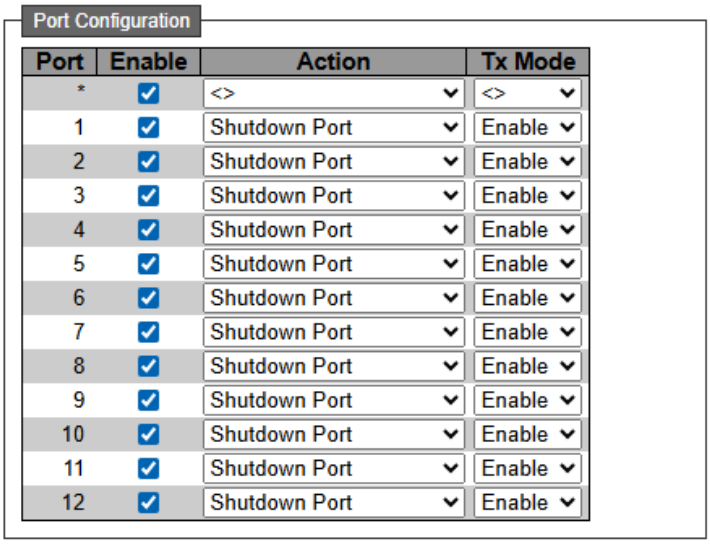

Parameter	Description
Port	The switch port number.
Enable	Controls whether loop protection is enabled on this switch port.
Action	Configures the action performed when a loop is detected on a port. Valid values are: **`Shutdown Port`**, **`Shutdown Port and Log`** or **`Log Only`**.
Tx Mode	Controls whether the port is actively generating loop protection PDU's, or whether it is just passively looking for looped PDU's.

4.16 Spanning Tree

Spanning Tree Protocol (STP) detects and breaks network loops and provides backup links between switches, bridges or routers. It allows a switch to interact with other STP compliant switches in your network to ensure that only one path exists between any two stations on the network.

There are three different versions of STP which are defined in different standards:

- IEEE 802.1D – **STP** (Spanning Tree Protocol)
- IEEE 802.1w – **RSTP** (Rapid Spanning Tree Protocol)
- IEEE 802.1s – **MSTP** (Multiple Spanning Tree Protocol)

Most widely supported and most commonly used is IEEE 802.1w RSTP (Rapid Spanning Tree Protocol) which allows faster convergence of the spanning tree than STP (while also being backwards compatible with STP-only aware bridges). In RSTP, topology change information is directly propagated throughout the network from the device that generates the topology change. In STP, a longer delay is required as the device that causes a topology change first notifies the root bridge and then the root bridge notifies the network. Multiple Spanning Tree Protocol (MSTP, IEEE 802.1s) extends RSTP by mapping multiple VLANs to different spanning tree instances. This allows traffic from different VLANs to use separate active paths, improving bandwidth utilization and network scalability while maintaining loop-free topologies.

Spanning Tree Terminology

- The **root bridge** is the base of the spanning tree. The bridge with the lowest priority value will be elected as root bridge. If multiple bridges have the lowest priority, the one of them with the lowest MAC address will be selected instead.

- Each bridge communicates with the root through the **root port**. The root port is the port on this switch with the lowest path cost to the root. If there is no root port, then this switch has been accepted as the root-bridge of the spanning tree network.
- **Path cost** is the cost of transmitting a frame onto a LAN through that port. The recommended cost is assigned according to the speed of the link to which a port is attached. The slower the media, the higher the cost.

	LINK SPEED	RECOMMENDED VALUE	RECOMMENDED RANGE	ALLOWED RANGE
Path Cost	4Mbps	250	100 to 1000	1 to 65535
Path Cost	10Mbps	100	50 to 600	1 to 65535
Path Cost	16Mbps	62	40 to 400	1 to 65535
Path Cost	100Mbps	19	10 to 60	1 to 65535
Path Cost	1Gbps	4	3 to 10	1 to 65535
Path Cost	10Gbps	2	1 to 5	1 to 65535

4.16.1 Bridge Settings

Spanning Tree status can be viewed under Monitor→Spanning Tree.

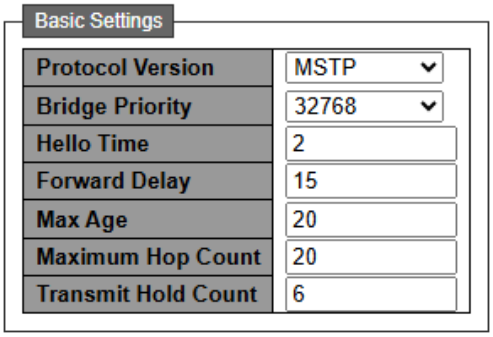

Parameter	Description
Protocol Version	The Spanning Tree protocol version setting. Valid values are: **MSTP**, **RSTP** and **STP**.
Bridge Priority	Controls the bridge priority. Lower numeric values have better priority. The bridge priority plus the MSTI instance number, concatenated with the 6-byte MAC address of the switch forms a *Bridge Identifier*. For **MSTP** operation, this is the priority of the CIST. Otherwise, this is the priority of the STP/RSTP bridge.
Hello Time	The interval between sending STP BPDU's. Valid values are in the range **1 to 10** seconds, default is 2 seconds. *Note: Changing this parameter from the default value is not recommended and may have adverse effects on your network.*
Max Age	The maximum age of the information transmitted by the Bridge when it is the Root Bridge. Valid values are in the range **6 to 40** seconds, and MaxAge must be <= (FwdDelay-1)*2.
Maximum Hop Count	This defines the initial value of remaining Hops for MSTI information generated at the boundary of an MSTI region. It defines how many bridges a root bridge can distribute its BPDU information to. Valid values are in the range **6 to 40** hops.
Transmit Hold Count	The number of BPDU's a bridge port can send per second. When exceeded, transmission of the next BPDU will be delayed. Valid values are in the range **1 to 10** BPDU's per second.

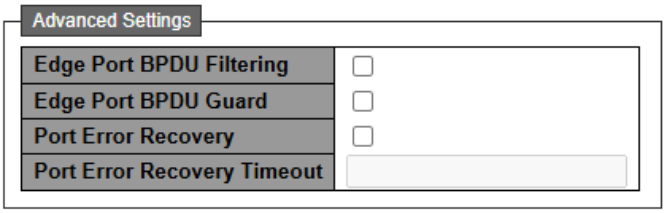

Parameter	Description
Edge Port BPDU Filtering	Control whether a port *explicitly* configured as **Edge** will transmit and receive BPDUs.
Edge Port BPDU Guard	Control whether a port *explicitly* configured as **Edge** will disable itself upon reception of a BPDU. The port will enter the *error-disabled* state, and will be removed from the active topology.
Port Error Recovery	Control whether a port in the *error-disabled* state automatically will be enabled after a certain time. If recovery is not enabled, ports have to be disabled and re-enabled for normal STP operation. The condition is also cleared by a system reboot.
Port Error Recovery Timeout	The time to pass before a port in the *error-disabled* state can be enabled. Valid values are between **`30 to 86400`** seconds (24 hours).

4.16.2 MSTI Mapping

An MSTI (Multiple Spanning Tree Instance) is a logical spanning tree within MSTP that maps one or more VLANs to a specific instance. Each MSTI maintains its own topology, port roles, and states, allowing different VLANs to use separate active paths and optimize traffic distribution while preventing loops.
Spanning Tree status can be viewed under Monitor→Spanning Tree.

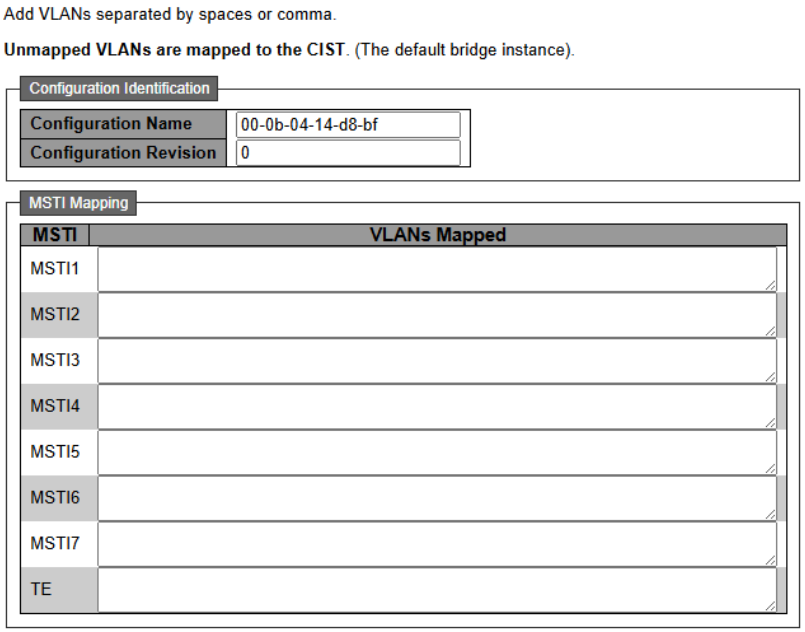

Parameter	Description
Configuration Name	The name identifying the VLAN to MSTI mapping. Bridges must share the name and revision (see below), as well as the VLAN-to-MSTI mapping configuration in order to share spanning trees for MSTI's (Intra-region). The name is at most 32 characters.
Configuration Revision	The revision of the MSTI configuration named above. This must be an integer between **`0`** **and** **`65535`**.
MSTI	The bridge instance. The CIST is not available for explicit mapping, as it will receive the VLANs not explicitly mapped.

	*Note: The **TE** (Traffic Engineering) instance is special, as it will not be controlled by MSTP itself. The TE instance will always be forwarding for all ports. The TE-MSTID is defined by IEEE 802.1Q-2018.*
VLANs Mapped	The list of VLANs mapped to the MSTI. The VLANs can be given as a single (**xx**, xx being between 1 and 4094) VLAN, or a range (**xx-yy**), each of which must be separated with comma and/or space. A VLAN can only be mapped to *one* MSTI. An unused MSTI should just be left empty. (I.e. not having any VLANs mapped to it.) Example: 2,5,20-40.

4.16.3 MSTI Priorities

Spanning Tree status can be viewed under Monitor→Spanning Tree.

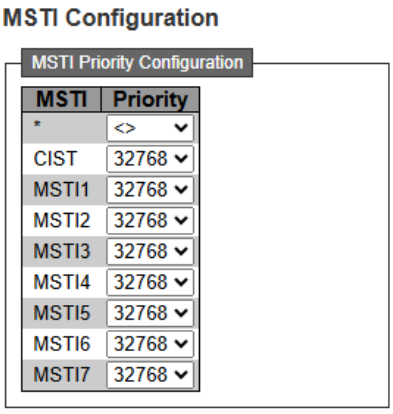

Parameter	Description
MSTI	The bridge instance. The CIST is the default instance, which is always active.
Priority	Controls the bridge priority. Lower numeric values have better priority. The bridge priority plus the MSTI instance number, concatenated with the 6-byte MAC address of the switch forms a *Bridge Identifier*.

4.16.4 CIST Ports

The CIST (Common and Internal Spanning Tree) is the single spanning tree in MSTP that connects all MST regions and legacy STP/RSTP bridges. It ensures overall loop-free connectivity across the network and coordinates topology changes between MSTIs and the external spanning tree environment.
Spanning Tree status can be viewed under Monitor→Spanning Tree.

STP CIST Port Configuration

CIST Aggregated Port Configuration

Port	STP Enabled	Path Cost	Priority	Admin Edge	Auto Edge	Restricted Role	Restricted TCN	BPDU Guard	Point-to-point
-	☑	Auto	128	Non-Edge	☑	☐	☐	☐	Forced True

CIST Normal Port Configuration

Port	STP Enabled	Path Cost	Priority	Admin Edge	Auto Edge	Restricted Role	Restricted TCN	BPDU Guard	Point-to-point
*	☑	<>	<>	<>	☑	☐	☐	☐	<>
1	☑	Auto	128	Non-Edge	☑	☐	☐	☐	Auto
2	☑	Auto	128	Non-Edge	☑	☐	☐	☐	Auto
3	☑	Auto	128	Non-Edge	☑	☐	☐	☐	Auto
4	☑	Auto	128	Non-Edge	☑	☐	☐	☐	Auto
5	☑	Auto	128	Non-Edge	☑	☐	☐	☐	Auto
6	☑	Auto	128	Non-Edge	☑	☐	☐	☐	Auto
7	☑	Auto	128	Non-Edge	☑	☐	☐	☐	Auto
8	☑	Auto	128	Non-Edge	☑	☐	☐	☐	Auto
9	☑	Auto	128	Non-Edge	☑	☐	☐	☐	Auto
10	☑	Auto	128	Non-Edge	☑	☐	☐	☐	Auto
11	☑	Auto	128	Non-Edge	☑	☐	☐	☐	Auto
12	☑	Auto	128	Non-Edge	☑	☐	☐	☐	Auto

Parameter	Description
Port	The switch port number of the logical STP port.
STP Enabled	Controls whether STP is enabled on this switch port.
Path Cost	Controls the path cost incurred by the port. The **`Auto`** setting will set the path cost as appropriate by the physical link speed, using the 802.1D recommended values. Using the **`Specific`** setting, a user-defined value can be entered. The path cost is used when establishing the active topology of the network. Lower path cost ports are chosen as forwarding ports in favour of higher path cost ports. Valid values are in the range **`1 to 200000000`**.
Priority	Controls the port priority. This can be used to control priority of ports having identical port cost. (See above). Lower priority is better.
operEdge (state flag)	Operational flag describing whether the port is connecting directly to edge devices. (*No* Bridges attached). Transition to the forwarding state is faster for edge ports (having *operEdge true*) than for other ports. The value of this flag is based on AdminEdge and AutoEdge fields. This flag is displayed as **Edge** in Monitor->Spanning Tree -> STP Detailed Bridge Status.
AdminEdge	Controls whether the *operEdge* flag should start as set or cleared. (The initial *operEdge* state when a port is initialized).
AutoEdge	Controls whether the bridge should enable automatic edge detection on the bridge port. This allows *operEdge* to be derived from whether BPDU's are received on the port or not.
Restricted Role	If enabled, causes the port not to be selected as Root Port for the *CIST* or any *MSTI*, even if it has the best spanning tree priority vector. Such a port will be selected as an Alternate Port after the Root Port has been selected. If set, it can cause lack of spanning tree connectivity. It can be set by a network administrator to prevent bridges external to a core region of the network influence the spanning tree active topology, possibly because those bridges are not under the full control of the administrator. This feature is also known as **`Root Guard`**.
Restricted TCN	If enabled, causes the port not to propagate received topology change notifications and topology changes to other ports. If set it can cause temporary loss of connectivity after changes in a spanning tree's active topology as a result of persistently incorrect learned station location information. It is set by a network administrator to prevent bridges external to a core region of the network, causing address flushing in that region, possibly because those bridges are not under the full control of the administrator or the physical link state of the attached LANs transits frequently.

BPDU Guard	If enabled, causes the port to disable itself upon receiving valid BPDU's. Contrary to the similar bridge setting, the port **Edge** status does not affect this setting. A port entering error-disabled state due to this setting is subject to the bridge Port Error Recovery setting as well.
Point-to-Point	Controls whether the port connects to a point-to-point LAN rather than to a shared medium. This can be automatically determined, or forced either true or false. Transition to the forwarding state is faster for point-to-point LANs than for shared media.

4.16.5 MSTI Ports

An MSTI port is a virtual port, which is instantiated separately for each active CIST (physical) port for each MSTI instance configured on and applicable to the port. The MSTI instance must be selected before displaying actual MSTI port configuration options. Select an MSTI and click *Get* to retrieve settings for the specific MSTI. Spanning Tree status can be viewed under Monitor-->Spanning Tree.

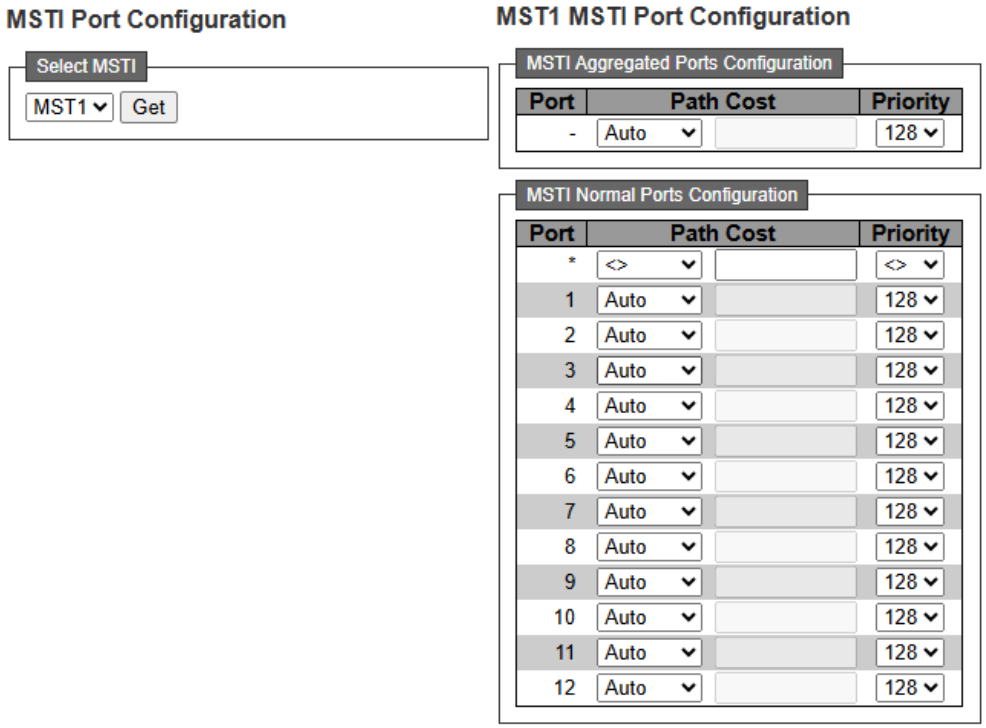

Parameter	Description
Port	The switch port number of the corresponding STP CIST (and MSTI) port.
Path Cost	Controls the path cost incurred by the port. The **Auto** setting will set the path cost as appropriate by the physical link speed, using the 802.1D recommended values. Using the **Specific** setting, a user-defined value can be entered. The path cost is used when establishing the active topology of the network. Lower path cost ports are chosen as forwarding ports in favor of higher path cost ports. Valid values are in the range **1 to 200000000**.
Priority	Controls the port priority. This can be used to control priority of ports having identical port cost. (See above). Lower priority is better.

4.17 IPMC Profile

IPMC (IP Multicast) Profile is used to deploy the access control on IP multicast streams. It is allowed to create at maximum 64 Profiles with at maximum 128 corresponding rules for each.

IPMC features (IGMP Snooping, MLD Snooping) can be configured under Configuration→IPMC.

4.17.1 Profile Table

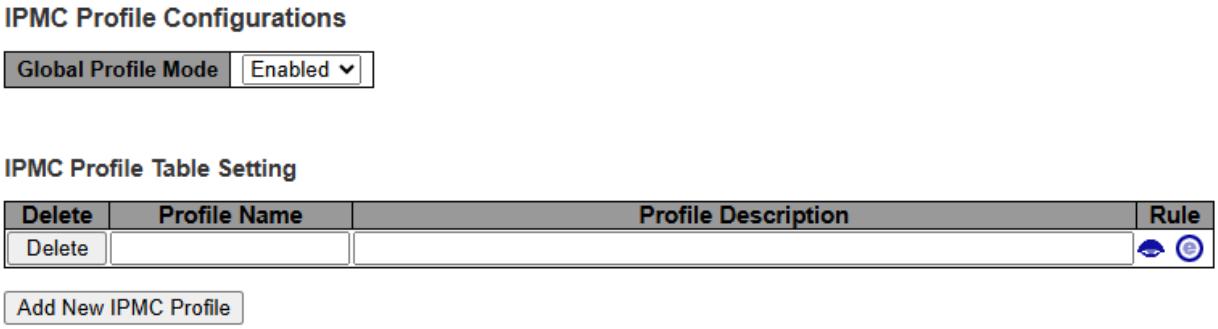

Parameter	Description
Global Profile Mode	Enable/Disable the Global IPMC Profile. System starts to do filtering based on profile settings only when the global profile mode is enabled.
Delete	Check to delete the entry. The designated entry will be deleted during the next save.
Profile Name	The name used for indexing the profile table. Each entry has the unique name which is composed of at maximum 16 alphabetic and numeric characters. At least one alphabet must be present.
Profile Description	Additional description, which is composed of at maximum 64 alphabetic and numeric characters, about the profile. No blank or space characters are permitted as part of description. Use "_" or "-" to separate the description sentence.
Rule	When the profile is created, click the edit button to enter the rule setting page of the designated profile. Summary about the designated profile will be shown by clicking the view button. You can manage or inspect the rules of the designated profile by using the following buttons: : View the rules associated with the designated profile. : Edit the rules associated with the designated profile.

When clicking on the or button, a new page will load showing the filtering rule settings for a specific IPMC profile. It displays the configured rule entries in precedence order. First rule entry has highest priority in lookup, while the last rule entry has lowest priority in lookup. Notice that the profile performs deny action for all groups if there is no any permit entry is included in the profile.

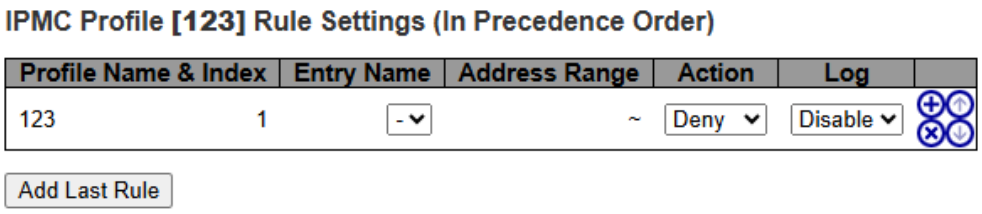

Parameter	Description
Profile Name	The name of the designated profile to be associated. This field is not editable.
Entry Name	The name used in specifying the address range used for this rule. Only the existing profile address entries will be chosen in the selected box. This field is not allowed to be selected as none ("-") while the Rule Settings Table is committed.
Address Range	The corresponding address range of the selected profile entry. This field is not editable and will be adjusted automatically according to the selected profile entry.
Action	Indicates the learning action upon receiving the Join/Report frame that has the group address matches the address range of the rule.

	Permit: Group address matches the range specified in the rule will be learned. **Deny**: Group address matches the range specified in the rule will be dropped.
Log	Indicates the logging preference upon receiving the Join/Report frame that has the group address matches the address range of the rule. **Enable**: Corresponding information of the group address, that matches the range specified in the rule, will be logged. **Disable**: Corresponding information of the group address, that matches the range specified in the rule, will not be logged.
Rule Management Buttons	You can manage rules and the corresponding precedence order by using the following buttons: : Insert a new rule before the current entry of rule. : Delete the current entry of rule. : Moves the current entry of rule up in the list. : Moves the current entry of rule down in the list.

4.17.2 Address Entry

The address entry is used to specify the address range that will be associated with IPMC Profile. It is allowed to create at maximum 128 address entries in the system.

IPMC Profile Address Configuration

Navigate Address Entry Setting in IPMC Profile by 20 entries per page.

Delete	Entry Name	Start Address	End Address
Delete			

Add New Address (Range) Entry

Parameter	Description
Delete	Check to delete the entry. The designated entry will be deleted during the next save.
Entry Name	The name used for indexing the address entry table. Each entry has the unique name which is composed of at maximum 16 alphabetic and numeric characters. At least one alphabet must be present.
Start Address	The starting IPv4/IPv6 Multicast Group Address that will be used as an address range.
End Address	The ending IPv4/IPv6 Multicast Group Address that will be used as an address range.

4.18 MVR

Multicast VLAN Registration (MVR) is a protocol for Layer 2 (IP)-networks that enables multicast-traffic from a source VLAN to be shared with subscriber-VLANs.
The main reason for using MVR is to save bandwidth by preventing duplicate multicast streams being sent in the core network. Instead, the stream(s) are received on the MVR-VLAN and forwarded to the VLANs where hosts have requested it/them.

In a multicast television application, a PC or a network television or a set-top box can receive the multicast stream. Multiple set-top boxes or PCs can be connected to one subscriber port, which is a switch port configured as an **MVR receiver port**. When a subscriber selects a channel, the set-top box or PC sends an

IGMP/MLD report message to switch A to join the appropriate multicast group address. Uplink ports that send and receive multicast data to and from the multicast VLAN are called **MVR source ports**.

The **Querier** should connect on the source port. By giving the static membership of MVR VLAN, the device only forwards the IGMP reports from downstream (receiver ports) to upstream (source ports) and the Query packet which comes from the downstream will be ignored silently.
After the MVR VLAN members are properly configured, it is required to associate an IPMC profile with the specific MVR VLAN to be the expected channel. The channel profile is defined by the IPMC Profile which provides the filtering conditions. Notice that the profile only works when the global profile mode is enabled. It is allowed to create at most 4 MVR VLANs with corresponding channel profile.

MVR statistics and information can be viewed at Monitor→MVR.

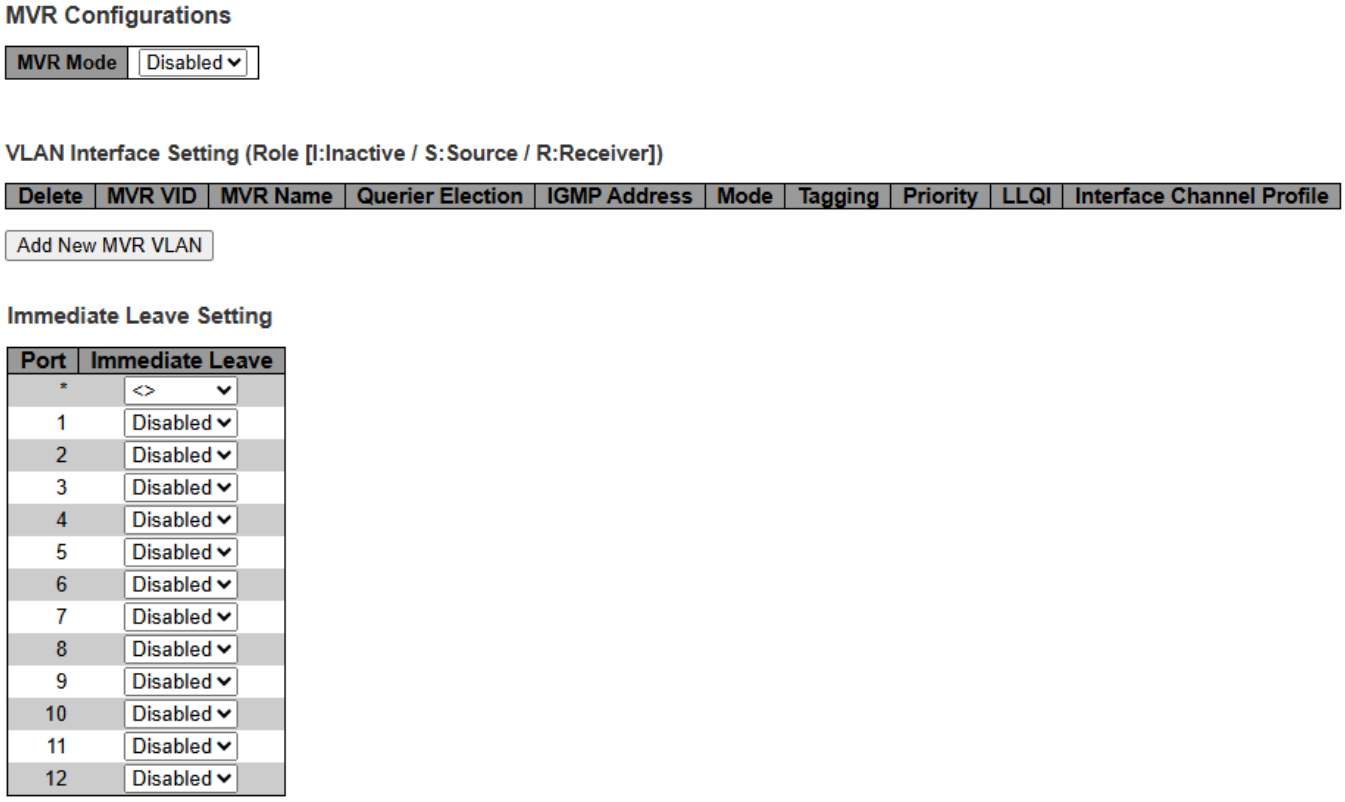

Parameter	Description
MVR Mode	Enable/Disable the Global MVR. The Unregistered Flooding control depends on the current configuration in IGMP/MLD Snooping. It is suggested to enable Unregistered Flooding control when the MVR group table is full.
Delete	Check to delete the entry. The designated entry will be deleted during the next save.
MVR VID	Specify the Multicast VLAN ID. *Caution: MVR source ports are not recommended to be overlapped with management VLAN ports.*
MVR Name	MVR Name is an optional attribute to indicate the name of the specific MVR VLAN. Maximum length of the MVR VLAN Name string is 16. MVR VLAN Name can only contain alphabets or numbers. When the optional MVR VLAN name is given, it should contain at least one alphabet. MVR VLAN name can be edited for the existing MVR VLAN entries or it can be added to the new entries.
Querier Election	Enable to join IGMP Querier election in the VLAN. Disable to act as an IGMP Non-Querier.
IGMP Address	The parameter is used when the device acts IGMP Querier. It defines the IPv4 address as source address used in IP header for IGMP control frames.

	The default IGMP address is not set (0.0.0.0). When the IGMP address is not set, system uses IPv4 management address of the IP interface associated with this VLAN. When the IPv4 management address is not set, system uses the first available IPv4 management address. Otherwise, system uses a pre-defined value. By default, this value will be 192.0.2.1.
Mode	Specify the MVR mode of operation. In Dynamic mode, MVR allows dynamic MVR membership reports on source ports. In Compatible mode, MVR membership reports are forbidden on source ports. The default is Dynamic mode.
Tagging	Specify whether the traversed IGMP/MLD control frames will be sent as Untagged or Tagged with MVR VID. The default is Tagged.
Priority	Specify how the traversed IGMP/MLD control frames will be sent in prioritized manner. The default Priority is 0.
LLQI	The parameter is used when the device acts as a Querier. It defines the maximum time to wait for IGMP/MLD report memberships on a receiver port before removing the port from multicast group membership. The value is in units of tenths of a second. The range is from **0 to 31744**. The default LLQI is 5 tenths or one-half second.
Interface Channel Profile	When the MVR VLAN is created, select the IPMC Profile as the channel filtering condition for the specific MVR VLAN. Summary about the Interface Channel Profiling (of the MVR VLAN) will be shown by clicking the view button. Profile selected for designated interface channel is not allowed to have overlapped permit group address. IPMC Profiles are managed at Configuration→IPMC Profile. : List the rules associated with the designated profile.
Port Role	Configure an MVR port of the designated MVR VLAN as one of the following roles. **Inactive**: The designated port does not participate MVR operations. **Source**: Configure uplink ports that receive and send multicast data as source ports. Subscribers cannot be directly connected to source ports. **Receiver**: Configure a port as a receiver port if it is a subscriber port and should only receive multicast data. It does not receive data unless it becomes a member of the multicast group by issuing IGMP/MLD messages. ***Caution**: MVR source ports are not recommended to be overlapped with management VLAN ports.* *Select the port role by clicking the Role symbol to switch the setting.* *"I" indicates Inactive; "S" indicates Source; "R" indicates Receiver. The default Role is Inactive.* Port 1 2 3 4 5 6 7 8 9 10 11 12 Role I S R I I I I I I I I I
Port	The logical port for the settings.
Immediate Leave	Enable the fast leave on the port. System will remove group record and stop forwarding data upon receiving the IGMPv2/MLDv1 leave message without sending last member query messages. It is recommended to enable this feature only when a single IGMPv2/MLDv1 host is connected to the specific port.

4.19 IPMC

IP Multicast (IPMC) enables efficient one-to-many traffic delivery by sending a single stream to multiple receivers. The switch manages multicast group membership and forwarding to ensure multicast traffic is delivered only to interested ports.

4.19.1 IGMP Snooping

IGMP snooping monitors IGMP messages between hosts and multicast routers to learn multicast group membership. The switch uses this information to forward multicast traffic only to interested ports, reducing unnecessary flooding and conserving bandwidth.

4.19.1.1 Basic Configuration

This contains the global and port configurations for IGMP Snooping.
IGMP Snooping status and information can be viewed at Monitor→IPMC→IGMP Snooping.

IGMP Snooping Configuration

Global Configuration	
Snooping Enabled	☑
Unregistered IPMCv4 Flooding Enabled	☑
IGMP SSM Range	232.0.0.0 / 8
Leave Proxy Enabled	☐
Proxy Enabled	☐

Parameter	Description
Snooping Enabled	Enable the Global IGMP Snooping.
Unregistered IPMCv4 Flooding Enabled	Enable unregistered IPMCv4 traffic flooding. The flooding control takes effect only when IGMP Snooping is enabled. When IGMP Snooping is disabled, unregistered IPMCv4 traffic flooding is always active in spite of this setting.
IGMP SSM Range	SSM (Source-Specific Multicast) Range allows the SSM-aware hosts and routers run the SSM service model for the groups in the address range. Requests made by IGMPv1 and IGMPv2 hosts in this range are ignored. Assign a valid IPv4 multicast address as prefix with a prefix length (from 4 to 32) for the range.
Leave Proxy Enabled	Enable IGMP Leave Proxy. This feature can be used to avoid forwarding unnecessary leave messages to the router side.
Proxy Enabled	Enable IGMP Proxy. This feature can be used to avoid forwarding unnecessary join and leave messages to the router side.

Port Related Configuration

Port	Router Port	Fast Leave	Throttling
*	☐	☐	<>
1	☐	☐	unlimited
2	☐	☐	unlimited
3	☐	☐	unlimited
4	☐	☐	unlimited
5	☐	☐	unlimited
6	☐	☐	unlimited
7	☐	☐	unlimited
8	☐	☐	unlimited
9	☐	☐	unlimited
10	☐	☐	unlimited
11	☐	☐	unlimited
12	☐	☐	unlimited

Parameter	Description
Port	The logical port for the settings.
Router Port	Specify which ports act as router ports. A router port is a port on the Ethernet switch that leads towards the Layer 3 multicast device or IGMP querier. If an aggregation member port is selected as a router port, the whole aggregation will act as a router port.
Fast Leave	Enable the fast leave on the port. System will remove group record and stop forwarding data upon receiving the IGMPv2 leave message without sending last member query messages. It is recommended to enable this feature only when a single IGMPv2 host is connected to the specific port.
Throttling	Enable to limit the number of multicast groups to which a switch port can belong.

4.19.1.2 VLAN Configuration

Each page shows up to 99 entries from the VLAN table, default being 20, selected through the "entries per page" input field. When first visited, the web page will show the first 20 entries from the beginning of the VLAN Table. The first displayed will be the one with the lowest VLAN ID found in the VLAN Table.
IGMP Snooping status and information can be viewed at Monitor→IPMC→IGMP Snooping.

For IGMP VLAN interface creation, you need to enter the IP configuration page to setup the IP interface first Configuration→System→IP→Add IP interface.

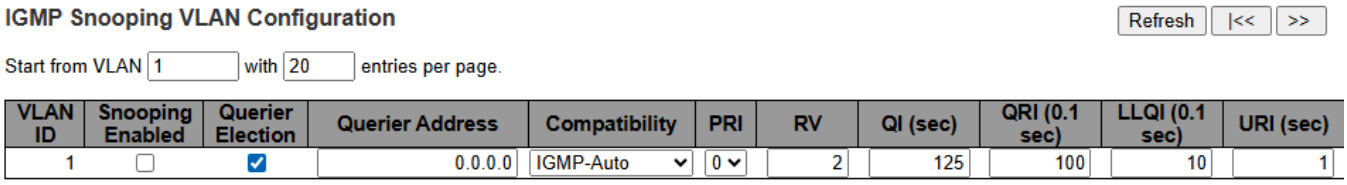

Parameter	Description
VLAN ID	The VLAN ID of the entry.
IGMP Snooping Enabled	Enable the per-VLAN IGMP Snooping. Up to 126 VLANs can be selected for IGMP Snooping.
Querier Election	Enable to join IGMP Querier election in the VLAN. Disable to act as an IGMP Non-Querier.
Querier Address	Define the IPv4 address as source address used in IP header for IGMP Querier election. When the Querier address is not set, system uses IPv4 management address of the IP interface associated with this VLAN. When the IPv4 management address is not set, system uses the first available IPv4 management address. Otherwise, system uses a pre-defined value. By default, this value will be 192.0.2.1.
Compatibility	Compatibility is maintained by hosts and routers taking appropriate actions depending on the versions of IGMP operating on hosts and routers within a network. The allowed selection is **IGMP-Auto**, **Forced IGMPv1**, **Forced IGMPv2**, **Forced IGMPv3**, default compatibility value is IGMP-Auto.
PRI	Priority of Interface. It indicates the IGMP control frame priority level generated by the system. These values can be used to prioritize different classes of traffic. The allowed range is **0 (best effort) to 7 (highest)**, default interface priority value is 0.
RV	Robustness Variable. The Robustness Variable allows tuning for the expected packet loss on a network. The allowed range is **2 to 255**, default robustness variable value is 2.
QI	Query Interval. The Query Interval is the interval between General Queries sent by the Querier. The allowed range is **1 to 31744** seconds, default query interval is 125 seconds.
QRI	Query Response Interval. The Maximum Response Delay used to calculate the Maximum Response Code inserted into the periodic General Queries. The allowed range is **0 to 31744** in tenths of seconds, default query response interval is 100 in tenths of seconds (10 seconds).
LLQI (LMQI for IGMP)	Last Member Query Interval. The Last Member Query Time is the time value represented by the Last Member Query Interval, multiplied by the Last Member Query Count. The allowed range is **0 to 31744** in tenths of seconds, default last member query interval is 10 in tenths of seconds (1 second).
URI	Unsolicited Report Interval. The Unsolicited Report Interval is the time between repetitions of a host's initial report of membership in a group. The allowed range is **0 to 31744** seconds, default unsolicited report interval is 1 second.

4.19.1.3 Port Filtering Profile

Assign an IPMC Profile to each port to enforce access control for multicast streams.
IGMP Snooping status and information can be viewed at Monitor→IPMC→IGMP Snooping.

IGMP Snooping Port Filtering Profile Configuration

Port	Filtering Profile
1	-
2	-
3	-
4	-
5	-
6	-
7	-
8	-
9	-
10	-
11	-
12	-

Parameter	Description
Port	The logical port for the settings.
Filtering Profile	Select the IPMC Profile as the filtering condition for the specific port. Summary about the designated profile will be shown by clicking the view button. : List the rules associated with the designated profile.

4.19.2 MLD Snooping

Multicast Listener Discovery (MLD) snooping performs for IPv6. What IGMP snooping does for IPv4: it monitors MLD messages to learn which ports are interested in specific multicast groups. Unlike IGMP, MLD operates over ICMPv6 instead of UDP, and it works with IPv6 addresses and link-local multicast behavior. The switch uses this information to forward IPv6 multicast traffic only to subscribed ports.

4.19.2.1 Basic Configuration

MLD Snooping status and information can be viewed at Monitor→IPMC→ MLD Snooping.

MLD Snooping Configuration

Global Configuration	
Snooping Enabled	☑
Unregistered IPMCv6 Flooding Enabled	☑
MLD SSM Range	ff3e:: / 96
Leave Proxy Enabled	☐
Proxy Enabled	☐

Parameter	Description
Snooping Enabled	Enable the Global IGMP Snooping.
Unregistered IPMCv6 Flooding Enabled	Enable unregistered IPMCv6 traffic flooding. The flooding control takes effect only when MLD Snooping is enabled. When MLD Snooping is disabled, unregistered IPMCv6 traffic flooding is always active in spite of this setting.
MLD SSM Range	SSM (Source-Specific Multicast) Range allows the SSM-aware hosts and routers run the SSM service model for the groups in the address range. Requests made by MLDv1 hosts in this range are ignored. Assign a valid IPv6 multicast address as prefix with a prefix length (from 8 to 128) for the range.
Leave Proxy Enabled	Enable MLD Leave Proxy. This feature can be used to avoid forwarding unnecessary leave messages to the router side.
Proxy Enabled	Enable MLD Proxy. This feature can be used to avoid forwarding unnecessary join and leave messages to the router side.

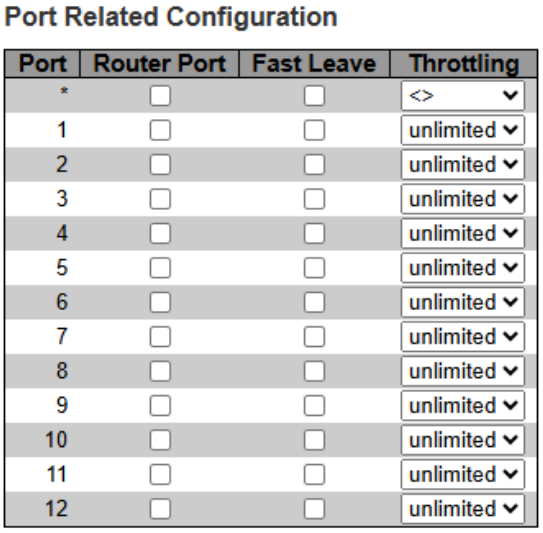

Parameter	Description
Port	The logical port for the settings.
Router Port	Specify which ports act as router ports. A router port is a port on the Ethernet switch that leads towards the Layer 3 multicast device or MLD querier. If an aggregation member port is selected as a router port, the whole aggregation will act as a router port.
Fast Leave	Enable the fast leave on the port. System will remove group record and stop forwarding data upon receiving the MLDv1 leave message without sending last member query messages. It is recommended to enable this feature only when a single MLDv1 host is connected to the specific port.
Throttling	Enable to limit the number of multicast groups to which a switch port can belong.

4.19.2.2 VLAN Configuration

MLD Snooping status and information can be viewed at Monitor→IPMC→MLD Snooping.

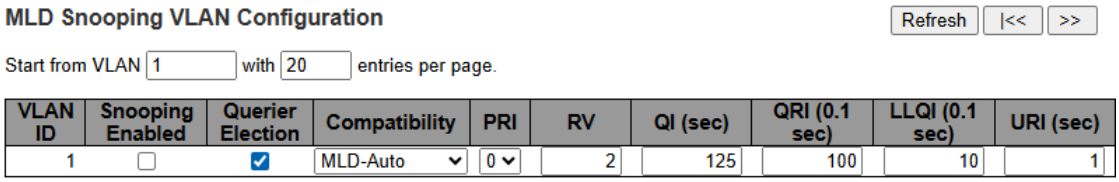

Parameter	Description
VLAN ID	The VLAN ID of the entry.
MLD Snooping Enabled	Enable the per-VLAN MLD Snooping. Up to 126 VLANs can be selected for MLD Snooping.
Querier Election	Enable to join MLD Querier election in the VLAN. Disable to act as a MLD Non-Querier.
Querier Address	Define the IPv4 address as source address used in IP header for IGMP Querier election. When the Querier address is not set, system uses IPv4 management address of the IP interface associated with this VLAN. When the IPv4 management address is not set, system uses the first available IPv4 management address. Otherwise, system uses a pre-defined value. By default, this value will be 192.0.2.1.
Compatibility	Compatibility is maintained by hosts and routers taking appropriate actions depending on the versions of MLD operating on hosts and routers within a network.

	The allowed selection is **IGMP-Auto**, **Forced MLDv1**, **Forced MLDv2**, default compatibility value is IGMP-Auto.
PRI	Priority of Interface. It indicates the MLD control frame priority level generated by the system. These values can be used to prioritize different classes of traffic. The allowed range is **0 (best effort) to 7 (highest)**, default interface priority value is 0.
RV	The Robustness Variable (RV) allows tuning for the expected packet loss on a network. The allowed range is **2 to 255**, default robustness variable value is 2.
QI	The Query Interval (QI) is the interval between General Queries sent by the Querier. The allowed range is **1 to 31744** seconds, default query interval is 125 seconds.
QRI	Query Response Interval. The Maximum Response Delay used to calculate the Maximum Response Code inserted into the periodic General Queries. The allowed range is **0 to 31744** in tenths of seconds, default query response interval is 100 in tenths of seconds (10 seconds).
LLQI	The Last Listener Query Interval (LLQI) is the Maximum Response Delay used to calculate the Maximum Response Code inserted into Multicast Address Specific Queries sent in response to Version 1 Multicast Listener Done messages. It is also the Maximum Response Delay used to calculate the Maximum Response Code inserted into Multicast Address and Source Specific Query messages. The allowed range is **0 to 31744** in tenths of seconds, default last listener query interval is 10 in tenths of seconds (1 second).
URI	The Unsolicited Report Interval (URI) is the time between repetitions of a host's initial report of membership in a group. The allowed range is **0 to 31744** seconds, default unsolicited report interval is 1 second.

4.19.2.3 Port Filtering Profile

MLD Snooping status and information can be viewed at Monitor→IPMC→ MLD Snooping.

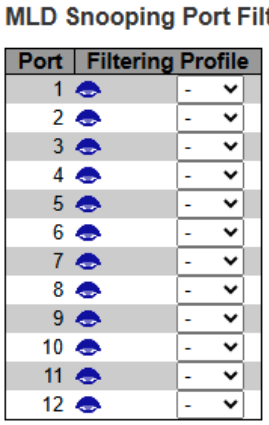

Parameter	Description
Port	The logical port for the settings.
Filtering Profile	Select the IPMC Profile as the filtering condition for the specific port. Summary about the designated profile will be shown by clicking the view button. : List the rules associated with the designated profile.

4.20 LLDP

The Link Layer Discovery Protocol (LLDP) specified in the IEEE 802.1ab standard allows stations attached to an IEEE 802 LAN to advertise device specific information and capabilities to other stations attached to the same IEEE 802 LAN. This includes the management address or addresses of the entity or entities that provide management of those capabilities, and the identification of the stations point of attachment to the IEEE 802 LAN required by those management entity or entities. The information distributed via this protocol is stored

by its recipients in a standard Management Information Base (MIB), making it possible for the information to be accessed by a Network Management System (NMS) using a management protocol such as the Simple Network Management Protocol (SNMP).

4.20.1 LLDP

Information about LLDP neighbors can be viewed at Monitor→LLDP.

LLDP Parameters

Tx Interval	30	seconds
Tx Hold	4	times
Tx Delay	2	seconds
Tx Reinit	2	seconds

Parameter	Description
Tx Interval	The switch periodically transmits LLDP frames to its neighbors for having the network discovery information up to date. The interval between each LLDP frame is determined by the *Tx Interval* value. Valid values are restricted to **5 - 32768** seconds.
Tx Hold	Each LLDP frame contains information about how long the information in the LLDP frame shall be considered valid. The LLDP information valid period is set to *Tx Hold* multiplied by *Tx Interval* seconds. Valid values are restricted to **2 - 10** times.
Tx Delay	If some configuration is changed (e.g. the IP address) a new LLDP frame is transmitted, but the time between the LLDP frames will always be at least the value of *Tx Delay* seconds. *Tx Delay* cannot be larger than 1/4 of the *Tx Interval* value. Valid values are restricted to **1 - 8192** seconds.
Tx Reinit	When an interface is disabled, LLDP is disabled or the switch is rebooted, a LLDP shutdown frame is transmitted to the neighboring units, signaling that the LLDP information isn't valid anymore. *Tx Reinit* controls the number of seconds between the shutdown frame and a new LLDP initialization. Valid values are restricted to **1 - 10** seconds.

TLV is an acronym for Type Length Value. A LLDP frame can contain multiple pieces of information. Each of these pieces of information is known as TLV.
For some TLVs it is configurable if the switch shall include the TLV in the LLDP frame. These TLVs are known as optional TLVs. If an optional TLVs is disabled, the corresponding information is not included in the LLDP frame.

LLDP Interface Configuration

				Optional TLVs				
Interface	Mode	CDP aware	Trap	Port Descr	Sys Name	Sys Descr	Sys Capa	Mgmt Addr
*	<>	☐	☐	☑	☑	☑	☑	☑
GigabitEthernet 1/1	Enabled	☐	☐	☑	☑	☑	☑	☑
GigabitEthernet 1/2	Enabled	☐	☐	☑	☑	☑	☑	☑
GigabitEthernet 1/3	Enabled	☐	☐	☑	☑	☑	☑	☑
GigabitEthernet 1/4	Enabled	☐	☐	☑	☑	☑	☑	☑
GigabitEthernet 1/5	Enabled	☐	☐	☑	☑	☑	☑	☑
GigabitEthernet 1/6	Enabled	☐	☐	☑	☑	☑	☑	☑
GigabitEthernet 1/7	Enabled	☐	☐	☑	☑	☑	☑	☑
GigabitEthernet 1/8	Enabled	☐	☐	☑	☑	☑	☑	☑
10GigabitEthernet 1/1	Enabled	☐	☐	☑	☑	☑	☑	☑
10GigabitEthernet 1/2	Enabled	☐	☐	☑	☑	☑	☑	☑
10GigabitEthernet 1/3	Enabled	☐	☐	☑	☑	☑	☑	☑
10GigabitEthernet 1/4	Enabled	☐	☐	☑	☑	☑	☑	☑

Parameter	Description
Interface	The switch interface name of the logical LLDP interface.
Mode	Select LLDP mode. **Rx only:** The switch will not send out LLDP information, but LLDP information from neighbor units is analyzed. **Tx only:** The switch will drop LLDP information received from neighbors but will send out LLDP information. **Disabled:** The switch will not send out LLDP information and will drop LLDP information received from neighbors. **Enabled:** The switch will send out LLDP information and will analyze LLDP information received from neighbors.
CDP Aware	Select CDP awareness. The CDP operation is restricted to decoding incoming CDP frames (The switch doesn't transmit CDP frames). CDP frames are only decoded if LLDP on the interface is enabled. Only CDP TLVs that can be mapped to a corresponding field in the LLDP neighbors' table are decoded. All other TLVs are discarded (Unrecognized CDP TLVs and discarded CDP frames are not shown in the LLDP statistics.). CDP TLVs are mapped onto LLDP neighbors' table as shown below. CDP TLV "Device ID" is mapped to the LLDP "Chassis ID" field. CDP TLV "Address" is mapped to the LLDP "Management Address" field. The CDP address TLV can contain multiple addresses, but only the first address is shown in the LLDP neighbors table. CDP TLV "Port ID" is mapped to the LLDP "Port ID" field. CDP TLV "Version and Platform" is mapped to the LLDP "System Description" field. Both the CDP and LLDP support "system capabilities", but the CDP capabilities cover capabilities that are not part of the LLDP. These capabilities are shown as "others" in the LLDP neighbors' table. If all interfaces have CDP awareness disabled the switch forwards CDP frames received from neighbor devices. If at least one interface has CDP awareness enabled all CDP frames are terminated by the switch. *Note: When CDP awareness on an interface is disabled the CDP information isn't removed immediately, but gets removed when the hold time is exceeded.*
Port Descr	Optional TLV: When checked the "port description" is included in LLDP information transmitted.
Sys Name	Optional TLV: When checked the "system name" is included in LLDP information transmitted.
Sys Descr	Optional TLV: When checked the "system description" is included in LLDP information transmitted.
Sys Capa	Optional TLV: When checked the "system capability" is included in LLDP information transmitted.
Mgmt Addr	Optional TLV: When checked the "management address" is included in LLDP information transmitted.

4.20.2 LLDP-MED

LLDP-MED is an extension of IEEE 802.1ab and is defined by the telecommunication industry association (TIA-1057). This function applies to VoIP devices which support LLDP-MED.
Information about LLDP neighbors can be viewed at Monitor→LLDP.

Fast Start Repeat Count

Rapid startup and Emergency Call Service Location Identification Discovery of endpoints is a critically important aspect of VoIP systems in general. In addition, it is best to advertise only those pieces of information which are specifically relevant to particular endpoint types (for example only advertise the voice network policy to permitted voice-capable devices), both in order to conserve the limited LLDPU space and to reduce security and system integrity issues that can come with inappropriate knowledge of the network policy.
With this in mind, LLDP-MED defines an LLDP-MED Fast Start interaction between the protocol and the application layers on top of the protocol, in order to achieve these related properties. Initially, a Network Connectivity Device will only transmit LLDP TLVs in an LLDPDU. Only after an LLDP-MED Endpoint Device is

detected, will an LLDP-MED capable Network Connectivity Device start to advertise LLDP-MED TLVs in outgoing LLDPDUs on the associated interface. The LLDP-MED application will temporarily speed up the transmission of the LLDPDU to start within a second, when a new LLDP-MED neighbor has been detected in order share LLDP-MED information as fast as possible to new neighbors.
Because there is a risk of an LLDP frame being lost during transmission between neighbors, it is recommended to repeat the fast start transmission multiple times to increase the possibility of the neighbors receiving the LLDP frame. With ***Fast start repeat count*** it is possible to specify the number of times the fast start transmission would be repeated. The recommended value is 4 times, given that 4 LLDP frames with a 1 second interval will be transmitted, when an LLDP frame with new information is received.
It should be noted that LLDP-MED and the LLDP-MED Fast Start mechanism is only intended to run on links between LLDP-MED Network Connectivity Devices and Endpoint Devices, and as such does not apply to links between LAN infrastructure elements, including Network Connectivity Devices, or other types of links.

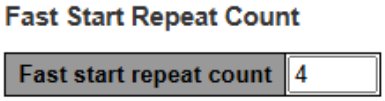

LLDP-MED Interface Configuration

It is possible to select which LLDP-MED information that shall be transmitted to the neighbors. When the checkbox is checked the information is included in the frame transmitted to the neighbor.

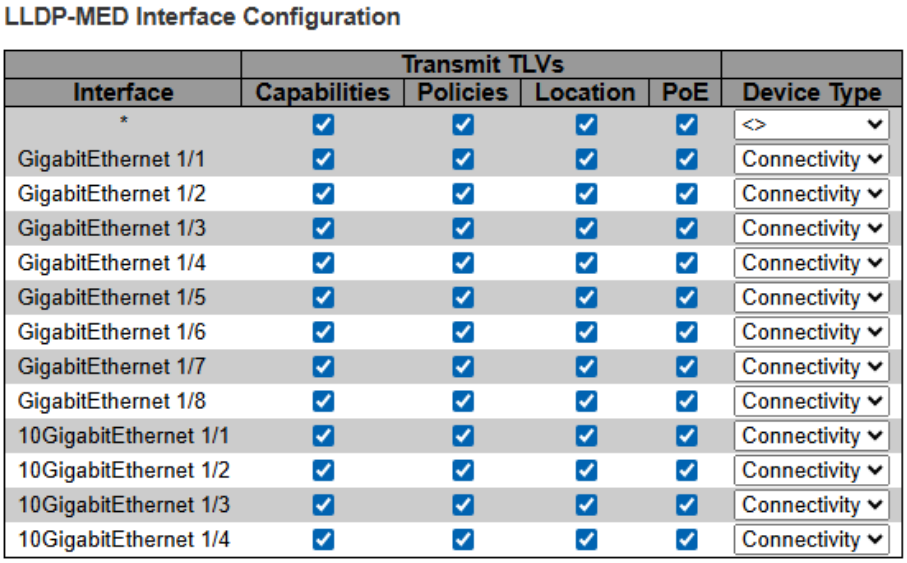

Parameter	Description
Interface	The interface name to which the configuration applies.
Capabilities	When checked, the switch's capabilities are included in LLDP-MED information transmitted.
Policies	When checked, the configured policies for the interface are included in LLDP-MED information transmitted.
Location	When checked the configured location information for the switch is included in LLDP-MED information transmitted.
PoE	When checked the configured PoE (Power Over Ethernet) information for the interface is included in LLDP-MED information transmitted.
Device Type	Any LLDP-MED Device is operating as a specific type of LLDP-MED Device, which may be either a Network **Connectivity** Device or a specific Class of **Endpoint** Device, as defined below. A Network **Connectivity** Device is an LLDP-MED Device that provides access to the IEEE 802 based LAN infrastructure for LLDP-MED **Endpoint** Devices

	An LLDP-MED Network **Connectivity** Device is a LAN access device based on any of the following technologies: 1. LAN Switch/Router 2. IEEE 802.1 Bridge 3. IEEE 802.3 Repeater (included for historical reasons) 4. IEEE 802.11 Wireless Access Point 5. Any device that supports the IEEE 802.1AB and MED extensions that can relay IEEE 802 frames via any method. An **Endpoint** Device is an LLDP-MED Device that sits at the network edge and provides some aspects of IP communications service, based on IEEE 802 LAN technology. The main difference between a Network **Connectivity** Device and an **Endpoint** Device is that only an **Endpoint** Device can start the LLDP-MED information exchange. Even though a switch always should be a Network **Connectivity** Device, it is possible to configure it to act as an **Endpoint** Device, and thereby start the LLDP-MED information exchange (In the case where two Network **Connectivity** Devices are connected together)

Coordinates Location

Parameter	Description
Latitude	**Latitude** SHOULD be normalized to within 0-90 degrees with a maximum of 4 digits. It is possible to specify the direction to either **North** of the equator or **South** of the equator.
Longitude	**Longitude** SHOULD be normalized to within 0-180 degrees with a maximum of 4 digits. It is possible to specify the direction to either **East** of the prime meridian or **West** of the prime meridian.
Altitude	**Altitude** SHOULD be normalized to within -2097151.9 to 2097151.9 with a maximum of 1 digits. It is possible to select between two altitude types (floors or meters). **Meters**: Representing meters of Altitude defined by the vertical datum specified. **Floors**: Representing altitude in a form more relevant in buildings which have different floor-to-floor dimensions. An altitude = 0.0 is meaningful even outside a building, and represents ground level at the given latitude and longitude. Inside a building, 0.0 represents the floor level associated with ground level at the main entrance.
Map Datum	The **Map Datum** is used for the coordinates given in these options: **WGS84**: (Geographical 3D) - World Geodesic System 1984, CRS Code 4327, Prime Meridian Name: Greenwich. **NAD83/NAVD88**: North American Datum 1983, CRS Code 4269, Prime Meridian Name: Greenwich; The associated vertical datum is the North American Vertical Datum of 1988 (NAVD88). This datum pair is to be used when referencing locations on land, not near tidal water (which would use Datum = NAD83/MLLW). **NAD83/MLLW**: North American Datum 1983, CRS Code 4269, Prime Meridian Name: Greenwich; The associated vertical datum is Mean Lower Low Water (MLLW). This datum pair is to be used when referencing locations on water/sea/ocean.

Civic Address Location

IETF Geopriv Civic Address based Location Configuration Information (Civic Address LCI). The total number of characters for the combined civic address information must not exceed 250 characters.

A couple of notes to the limitation of 250 characters.

1) If more than one civic address location is used, each of the additional civic address locations will use 2 extra characters in addition to the civic address location text.
2) The 2 letter country code is not part of the 250 characters limitation.

Civic Address Location

Country code		State		County	
City		City district		Block (Neighborhood)	
Street		Leading street direction		Trailing street suffix	
Street suffix		House no.		House no. suffix	
Landmark		Additional location info		Name	
Zip code		Building		Apartment	
Floor		Room no.		Place type	
Postal community name		P.O. Box		Additional code	

Parameter	Description
Country code	The two-letter ISO 3166 country code in capital ASCII letters - Example: DK, DE or US.
State	National subdivisions (state, canton, region, province, prefecture).
Altitude	County, parish, gun (Japan), district.
City	City, township, shi (Japan) - Example: Copenhagen.
City district	City division, borough, city district, ward, chou (Japan).
Block (Neighborhood)	Neighborhood, block.
Street	Street - Example: Poppelvej.
Leading street direction	Leading street direction - Example: N.
Trailing street suffix	Trailing street suffix - Example: SW.
Street suffix	Street suffix - Example: Ave, Platz.
House no.	House number suffix - Example: A, 1/2.
Landmark	Landmark or vanity address - Example: Columbia University.
Additional location info	Additional location info - Example: South Wing.
Name	Name (residence and office occupant) - Example: Flemming Jahn.
Zip code	Postal/zip code - Example: 2791.
Building	Building (structure) - Example: Low Library.
Floor	Floor - Example: 4.
Room no.	Room number - Example: 450F.
Place type	Place type - Example: Office.
Postal community name	Postal community name - Example: Leonia.
P.O. Box	Post office box (P.O. BOX) - Example: 12345.
Additional code	Additional code - Example: 1320300003.

Emergency Call Services

Emergency Call Service (e.g. E911 and others), such as defined by TIA or NENA. Emergency Call Service ELIN identifier data format is defined to carry the ELIN identifier as used during emergency call setup to a traditional CAMA or ISDN trunk-based PSAP. This format consists of a numerical digit string, corresponding to the ELIN to be used for emergency calling.

Emergency Call Service	

Policies

Network Policy Discovery enables the efficient discovery and diagnosis of mismatch issues with the VLAN configuration, along with the associated Layer 2 and Layer 3 attributes, which apply for a set of specific protocol applications on that port. Improper network policy configurations are a very significant issue in VoIP environments that frequently result in voice quality degradation or loss of service.
Policies are only intended for use with applications that have specific 'real-time' network policy requirements, such as interactive voice and/or video services.

The network policy attributes advertised are:

1. Layer 2 VLAN ID (IEEE 802.1Q-2003)
2. Layer 2 priority value (IEEE 802.1D-2004)
3. Layer 3 Diffserv code point (DSCP) value (IETF RFC 2474)

This network policy is potentially advertised and associated with multiple sets of application types supported on a given port. The application types specifically addressed are:

1. Voice
2. Guest Voice
3. Softphone Voice
4. Video Conferencing
5. Streaming Video
6. Control / Signaling (conditionally support a separate network policy for the media types above)

A large network may support multiple VoIP policies across the entire organization, and different policies per application type. LLDP-MED allows multiple policies to be advertised per port, each corresponding to a different application type. Different ports on the same Network Connectivity Device may advertise different sets of policies, based on the authenticated user identity or port configuration.

It should be noted that LLDP-MED is not intended to run on links other than between Network Connectivity Devices and Endpoints, and therefore does not need to advertise the multitude of network policies that frequently run on an aggregated link interior to the LAN.

Policies

Delete	Policy ID	Application Type	Tag	VLAN ID	L2 Priority	DSCP
Delete	0	Voice	Tagged	1	0	0

Parameter	Description
Delete	Check to delete the policy. It will be deleted during the next save.
Policy ID	ID for the policy. This is auto generated and shall be used when selecting the policies that shall be mapped to the specific interfaces.
Application Type	Intended use of the application types: **Voice** - for use by dedicated IP Telephony handsets and other similar appliances supporting interactive voice services. These devices are typically deployed on a separate VLAN for ease of deployment and enhanced security by isolation from data applications.

	Voice Signaling (conditional) - for use in network topologies that require a different policy for the voice signaling than for the voice media. This application type should not be advertised if all the same network policies apply as those advertised in the **Voice** application policy. **Guest Voice** - support a separate 'limited feature-set' voice service for guest users and visitors with their own IP Telephony handsets and other similar appliances supporting interactive voice services. **Guest Voice Signaling** (conditional) - for use in network topologies requiring a different policy for guest voice signaling than for the guest voice media. This application type should not be advertised if all the same network policies apply as those advertised in the **Guest Voice** application policy. **Softphone Voice** - for use by softphone applications on typical data centric devices, such as PCs or laptops. This class of endpoints frequently does not support multiple VLANs, if at all, and are typically configured to use an 'untagged' VLAN or a single 'tagged' data specific VLAN. When a network policy is defined for use with an 'untagged' VLAN (see Tagged flag below), then the L2 priority field is ignored and only the DSCP value has relevance. **Video Conferencing** - for use by dedicated Video Conferencing equipment and other similar appliances supporting real-time interactive video/audio services. **Streaming Video** - for use by broadcast or multicast-based video content distribution and other similar applications supporting streaming video services that require specific network policy treatment. Video applications relying on TCP with buffering would not be an intended use of this application type. **Video Signaling** (conditional) - for use in network topologies that require a separate policy for the video signaling than for the video media. This application type should not be advertised if all the same network policies apply as those advertised in the **Video Conferencing** application policy.
Tag	*Tag* indicating whether the specified application type is using a 'tagged' or an 'untagged' VLAN. **Untagged** indicates that the device is using an untagged frame format and as such does not include a tag header as defined by IEEE 802.1Q-2003. In this case, both the VLAN ID and the Layer 2 priority fields are ignored and only the DSCP value has relevance. **Tagged** indicates that the device is using the IEEE 802.1Q tagged frame format, and that both the VLAN ID and the Layer 2 priority values are being used, as well as the DSCP value. The tagged format includes an additional field, known as the tag header. The tagged frame format also includes priority tagged frames as defined by IEEE 802.1Q-2003.
VLAN ID	VLAN identifier (VID) for the interface as defined in IEEE 802.1Q-2003.
L2 Priority	*L2 Priority* is the Layer 2 priority to be used for the specified application type. *L2 Priority* may specify one of eight priority levels (0 through 7), as defined by IEEE 802.1D-2004. A value of 0 represents use of the default priority as defined in IEEE 802.1D-2004.
DSCP	*DSCP* value to be used to provide Diffserv node behavior for the specified application type as defined in IETF RFC 2474. *DSCP* may contain one of 64 code point values (0 through 63). A value of 0 represents use of the default DSCP value as defined in RFC 2475.

Policy Interface Configuration

Every interface may advertise a unique set of network policies or different attributes for the same network policies, based on the authenticated user identity or interface configuration. The number of policies supported is 32.

Policy Interface Configuration

Interface	0	1	2
GigabitEthernet 1/1	☐	☐	☐
GigabitEthernet 1/2	☐	☐	☐
GigabitEthernet 1/3	☐	☐	☐
GigabitEthernet 1/4	☐	☐	☐
GigabitEthernet 1/5	☐	☐	☐
GigabitEthernet 1/6	☐	☐	☐
GigabitEthernet 1/7	☐	☐	☐
GigabitEthernet 1/8	☐	☐	☐
10GigabitEthernet 1/1	☐	☐	☐
10GigabitEthernet 1/2	☐	☐	☐
10GigabitEthernet 1/3	☐	☐	☐
10GigabitEthernet 1/4	☐	☐	☐

Parameter	Description
Interface	The interface name to which the configuration applies.
Policy ID	The set of policies that shall apply to a given interface. The set of policies is selected by check marking the checkboxes that correspond to the policies.

4.21 PoE

Power budget note for 9561-8GP4XS-TSN: the PoE System Configuration page accepts a Power supply value from 0 to 480 W. The PoE delivery budget available to powered devices is 360 W.

Power Over Ethernet (PoE) is used to transmit electrical power, to remote devices over standard Ethernet cable. It could for example be used for powering IP cameras, IP telephones, wireless LAN access points and other equipment, where it would be difficult or expensive to connect the equipment to main power supply.

PoE status and usage information can be viewed at Monitor→PoE.

System Configuration

Power supply	400 W (max 480 W) poe power limit: 360 W
Interruptible power	Disabled
PD Auto-class request	Disabled
Legacy PD-Class Mode	Standard

Parameter	Description
Power supply	Specifies the available PoE power supply value in Watts. The configurable range shown by the Web GUI is 0 - 480 W.
poe power limit	Shows the PoE delivery budget available to powered devices. For 9561-8GP4XS-TSN, the PoE power limit is 360 W.
Interruptible power	Controls if PoE power should be interrupted (shut down for 5 Sec) during unit software restart cycle or remain unchanged during the unit software restart cycle. **Disabled**: PoE power remains unchanged during the entire unit software restart cycle. **Enabled**: PoE power to already powered PD devices will be turned off for 5 Sec during unit software restart cycle.
PD Auto-class request	*Controls whether PD Auto-class request is used. When enabled, a PD that supports IEEE 802.3bt Auto-class can report its effective maximum power consumption to the PSE so the switch can allocate power based on the measured/requested demand instead of only the configured port Type.*

Legacy PD-Class Mode	Legacy PD-Class Mode (Standard / PoH / Ignore Pd-Class) applies only to PoE ports configured as Plus mode. Ignore Pd-Class applies only to ports configured as Type3-60W or Type4-90W; Type3-15W and Type3-30W operate as Standard mode.

Note:

- Legacy PD-Class Mode (standard / PoH / Ignore Pd-Class) configuration is applicable only to PoE ports configured as 'Plus' under the mode column.

- Legacy PD-Class Mode - 'Ignore Pd-Class' is applicable only for ports configured as Type3-60W, Type4-90W. Excluding Type3-15W, Type3-30W.

Port Configuration

Port	Type	Mode	Pwr Mng	Priority	LLDP	Cable Length
*	<>	<>	<>	<>	<>	<>
1	type4-90w	plus	dynamic	low	disable	max60
2	type4-90w	plus	dynamic	low	enable	max100
3	type4-90w	plus	dynamic	low	enable	max100
4	type4-90w	plus	dynamic	low	disable	max60
5	type4-90w	plus	dynamic	low	enable	max100
6	type4-90w	plus	dynamic	low	enable	max100
7	type4-90w	plus	dynamic	low	enable	max100
8	type4-90w	plus	dynamic	low	enable	max100

Parameter	Description
Port	Switch port number. Only PoE-capable ports are shown.
Type	Configures the maximum PoE power type for the port. Supported values are Type3-15W, Type3-30W, Type3-60W, and Type4-90W. Power demotion may allocate a lower power level when the PD request or available system budget requires it. Type3-15W limits the port to 15 W, Type3-30W limits the port to 30 W, Type3-60W limits the port to 60 W/30 W for four-pair/two-pair operation, and Type4-90W limits the port to 90 W/45 W for four-pair/two-pair operation.
Mode	Configure PoE port to one of the following options. **Disabled**: PoE port is Disabled. Port becomes a non-PoE switch Ethernet port. **Standard**: PoE port is enabled and compliant with IEEE 802.3bt specification. **Plus**: PoE port is enabled and supports non IEEE-802.3-af/at/bt PoE PDs.
Power Management Mode	Configures how free PoE power is calculated for additional PDs. Dynamic deducts actual PD power consumption. Static deducts allocated/requested power according to the configured port Type and negotiated PD request. Hybrid behaves dynamically unless a maximum PoE power is negotiated over LLDP, then limits the port to the negotiated power.
Priority	PoE port priority controls the order of the PoE ports during Power-On sequence and during Power-Off sequence whenever overall power consumption exceeds the maximum available power. All ports configured as Critical are turned on first, followed by all ports configured as High, and lastly all ports configured as Low. For all ports of the same priority, the lowest PoE switch port number will be turned on first, followed by the next highest PoE switch port number. For example, if all ports are set to priority-Low, then port #1 will be turned on first, followed by port #2, etc. The highest PoE switch port number will be the first port to be turned Off whenever the overall power consumption exceeds the maximum available power, again in accordance with PoE port number, while giving priority to Critical over High and High over Low. **Critical**: Highest level PoE port priority **High**: Mid-level PoE port priority **Low**: Lowest level PoE port priority
LLDP	The LLDP configures the PoE port behavior with respect to LLDP packets from PD. **Enable**: PoE parameters as PD power-request received through LLDP are processed. **Disable**: PoE parameters received through LLDP are ignored. LLDP protocol is configured by its own configuration web page and transmission of PoE information via LLDP can be configured with the LLDP-MED configuration page.

Cable Length	Cable length assists port power optimization allocation to remote PD advertising their power requirement over LLDP. Cable power loss for PD Type4 requesting 71.3W is assumed to be 18.7W for 100m cable length. So the port max power will be set to 90W. However, if PD is located only 30 meters away then port maximum power can be lowered to 76W freeing 14W for other PDs. **max10**: Ethernet cable length is 10 meter or less. **max30**: Ethernet cable length is 30 meter or less. **max60**: Ethernet cable length is 60 meter or less. **max100**: Ethernet cable length is 100 meter or less.

Figure: PoE configuration page showing system power budget fields and per-port PoE settings.

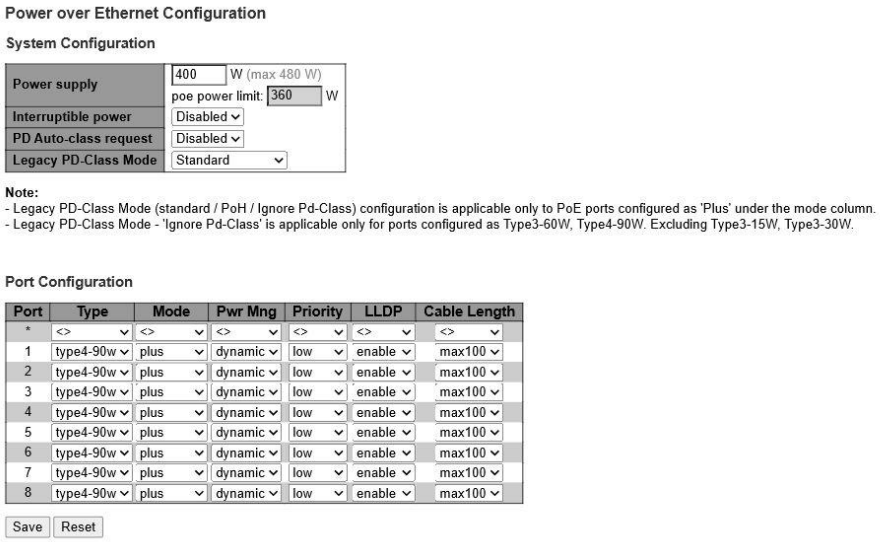

4.22 MAC Table

Switching of frames is based upon the DMAC address contained in the frame. The switch builds up a table that maps MAC addresses to switch ports for knowing which ports the frames should go to (based upon the DMAC address in the frame). This table contains both static and dynamic entries. The static entries are configured by the network administrator if the administrator wants to do a fixed mapping between the DMAC address and switch ports.

The frames also contain a source MAC address (SMAC address), which shows the MAC address of the equipment sending the frame. The SMAC address is used by the switch to automatically update the MAC table with these dynamic MAC addresses. Dynamic entries are removed from the MAC table if no frame with the corresponding SMAC address has been seen after a configurable age time.

Set timeouts for entries in the dynamic MAC Table and configure the static MAC table here.
The MAC Table entries can be viewed at Monitor→MAC Table.

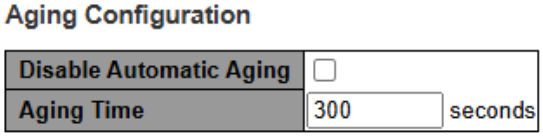

By default, dynamic entries are removed from the MAC table after 300 seconds. This removal is also called aging.

Configure aging time by entering a value here in seconds. The allowed range is **10** to **1000000** seconds.
Disable the automatic aging of dynamic entries by checking **Disable automatic aging**.

MAC Table Learning

	Port Members											
	1	2	3	4	5	6	7	8	9	10	11	12
Auto	◉	◉	◉	◉	◉	◉	◉	◉	◉	◉	◉	◉
Disable	○	○	○	○	○	○	○	○	○	○	○	○
Secure	○	○	○	○	○	○	○	○	○	○	○	○

If the learning mode for a given port is greyed out, another module is in control of the mode, so that it cannot be changed by the user. An example of such a module is the MAC-Based Authentication under 802.1X.
Each port can do learning based upon the following settings:

Parameter	Description
Auto	Learning is done automatically as soon as a frame with unknown SMAC is received.
Disable	No learning is done.
Secure	Only static MAC entries are learned, all other frames are dropped. *Note: Make sure that the link used for managing the switch is added to the Static Mac Table before changing to secure learning mode, otherwise the management link is lost and can only be restored by using another non-secure port or by connecting to the switch via the serial interface.*

VLAN Learning Configuration

Learning-disabled VLANs	

This field shows the Learning-disabled VLANs. When a new MAC arrives into a learning-disabled VLAN, the MAC won't be learnt. By the default, the field is empty. More VLANs may be created by using a list syntax where the individual elements are separated by commas. Ranges are specified with a dash separating the lower and upper bound.
The following example will create VLANs 1, 10, 11, 12, 13, 200, and 300: 1,10-13,200,300. Spaces are allowed in between the delimiters.

Static MAC Table Configuration

			Port Members											
Delete	VLAN ID	MAC Address	1	2	3	4	5	6	7	8	9	10	11	12
Delete	1	00-00-00-00-00-00	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐

The static entries in the MAC table are shown in this table. The static MAC table can contain 64 entries.
The MAC table is sorted first by VLAN ID and then by MAC address.

Parameter	Description
Delete	Learning is done automatically as soon as a frame with unknown SMAC is received.
VLAN ID	No learning is done.
MAC Address	Only static MAC entries are learned, all other frames are dropped. *Note: Make sure that the link used for managing the switch is added to the Static Mac Table before changing to secure learning mode, otherwise the management link is lost and can only be restored by using another non-secure port or by connecting to the switch via the serial interface.*
Port Members	Checkmarks indicate which ports are members of the entry. Check or uncheck as needed to modify the entry.

4.23 VLANs

Virtual LAN (VLAN) is a method to restrict communication between switch ports. A VLAN logically segments a physical network at layer 2 into separate broadcast domains. It improves traffic isolation, security, and manageability by grouping devices independently of their physical location.

4.23.1 Configuration

This page allows for controlling VLAN configuration on the switch.
The page is divided into a global section and a per-port configuration section.
VLAN port settings and status can be viewed at Monitor→VLANs.

Global VLAN Configuration

Allowed Access VLANs	1
Ethertype for Custom S-ports	88A8

Parameter	Description
Allowed Access VLANs	This field shows the allowed Access VLANs, i.e. it only affects ports configured as Access ports. Ports in other modes are members of the VLANs specified in the Allowed VLANs field. By default, only VLAN 1 is enabled. More VLANs may be created by using a list syntax where the individual elements are separated by commas. Ranges are specified with a dash separating the lower and upper bound. The following example will create VLANs 1, 10, 11, 12, 13, 200, and 300: 1,10-13,200,300. Spaces are allowed in between the delimiters.
Ethertype for Custom S-ports	This field specifies the ethertype/TPID (specified in hexadecimal) used for Custom S-ports. The setting is in force for all ports whose Port Type is set to S-Custom-Port.

Port VLAN Configuration

Port	Mode	Port VLAN	Port Type	Ingress Filtering	Ingress Acceptance	Egress Tagging	Allowed VLANs	Forbidden VLANs
*	<>	1	<>	☑	<>	<>	1	
1	Access	1	C-Port	☑	Tagged and Untagged	Untag All	1	
2	Access	1	C-Port	☑	Tagged and Untagged	Untag All	1	
3	Access	1	C-Port	☑	Tagged and Untagged	Untag All	1	
4	Access	1	C-Port	☑	Tagged and Untagged	Untag All	1	
5	Access	1	C-Port	☑	Tagged and Untagged	Untag All	1	
6	Access	1	C-Port	☑	Tagged and Untagged	Untag All	1	
7	Access	1	C-Port	☑	Tagged and Untagged	Untag All	1	
8	Access	1	C-Port	☑	Tagged and Untagged	Untag All	1	
9	Access	1	C-Port	☑	Tagged and Untagged	Untag All	1	
10	Access	1	C-Port	☑	Tagged and Untagged	Untag All	1	
11	Access	1	C-Port	☑	Tagged and Untagged	Untag All	1	
12	Access	1	C-Port	☑	Tagged and Untagged	Untag All	1	

Parameter	Description
Port	This is the logical port number of this row.
Mode	The port mode (default is Access) determines the fundamental behavior of the port in question. A port can be in one of three modes as described below. Whenever a particular mode is selected, the remaining fields in that row will be either grayed out or made changeable depending on the mode in question. Grayed out fields show the value that the port will get when the mode is applied. **Access:** Access ports are normally used to connect to end stations. Dynamic features like Voice VLAN may add the port to more VLANs behind the scenes. Access ports have the following characteristics: • Member of exactly one VLAN, the Port VLAN (a.k.a. Access VLAN), which by default is 1

	• Accepts untagged and C-tagged frames • Discards all frames not classified to the Access VLAN • On egress all frames are transmitted untagged **Trunk:** Trunk ports can carry traffic on multiple VLANs simultaneously and are normally used to connect to other switches. Trunk ports have the following characteristics: • By default, a trunk port is member of all VLANs (1-4095) • The VLANs that a trunk port is member of may be limited by the use of *Allowed VLANs* • Frames classified to a VLAN that the port is not a member of are discarded • By default, all frames but frames classified to the Port VLAN (a.k.a. Native VLAN) get tagged on egress. Frames classified to the Port VLAN do not get C-tagged on egress • Egress tagging can be changed to tag all frames, in which case only tagged frames are accepted on ingress **Hybrid:** Hybrid ports resemble trunk ports in many ways, but add additional port configuration features. In addition to the characteristics described for trunk ports, hybrid ports have these abilities: • Can be configured to be VLAN tag unaware, C-tag aware, S-tag aware, or S-custom-tag aware • Ingress filtering can be controlled • Ingress acceptance of frames and configuration of egress tagging can be configured independently
Port VLAN	Determines the port's VLAN ID (a.k.a. PVID). Allowed VLANs are in the range 1 through 4095, default being 1. On ingress, frames get classified to the Port VLAN if the port is configured as VLAN unaware, the frame is untagged, or VLAN awareness is enabled on the port, but the frame is priority tagged (VLAN ID = 0). On egress, frames classified to the Port VLAN do not get tagged if Egress Tagging configuration is set to untag Port VLAN. The Port VLAN is called an "Access VLAN" for ports in Access mode and Native VLAN for ports in Trunk or Hybrid mode.
Port Type	Ports in hybrid mode allow for changing the port type, that is, whether a frame's VLAN tag is used to classify the frame on ingress to a particular VLAN, and if so, which TPID it reacts on. Likewise, on egress, the Port Type determines the TPID of the tag, if a tag is required. **Unaware:** On ingress, all frames, whether carrying a VLAN tag or not, get classified to the Port VLAN, and possible tags are not removed on egress. **C-Port:** On ingress, frames with a VLAN tag with TPID = 0x8100 get classified to the VLAN ID embedded in the tag. If a frame is untagged or priority tagged, the frame gets classified to the Port VLAN. If frames must be tagged on egress, they will be tagged with a C-tag. **S-Port:** On egress, if frames must be tagged, they will be tagged with an S-tag. On ingress, frames with a VLAN tag with TPID = 0x88A8 get classified to the VLAN ID embedded in the tag. Priority-tagged frames are classified to the Port VLAN. If the port is configured to accept **Tagged Only** frames (see *Ingress Acceptance* below), frames without this TPID are dropped. **S-Custom-Port:** On egress, if frames must be tagged, they will be tagged with the custom S-tag. On ingress, frames with a VLAN tag with a TPID equal to the *Ethertype configured for Custom-S ports* get classified to the VLAN ID embedded in the tag. Priority-tagged frames are classified to the Port VLAN. If the port is configured to accept **Tagged Only** frames (see *Ingress Acceptance* below), frames without this TPID are dropped.
Ingress Filtering	Hybrid ports allow for changing ingress filtering. Access and Trunk ports always have ingress filtering enabled. If ingress filtering is enabled (checkbox is checked), frames classified to a VLAN that the port is not a member of get discarded. If ingress filtering is disabled, frames classified to a VLAN that the port is not a member of are accepted and forwarded to the switch engine. However, the port will never transmit frames classified to VLANs that it is not a member of.

Ingress Acceptance	Hybrid ports allow for changing the type of frames that are accepted on ingress. **Tagged and Untagged** Both tagged and untagged frames are accepted. See *Port Type* for a description of when a frame is considered tagged. **Tagged Only** Only frames tagged with the corresponding *Port Type* tag are accepted on ingress. **Untagged Only** Only untagged frames are accepted on ingress. See *Port Type* for a description of when a frame is considered untagged.
Egress Tagging	Ports in Trunk and Hybrid mode may control the tagging of frames on egress. **Untag Port VLAN** Frames classified to the Port VLAN are transmitted untagged. Other frames are transmitted with the relevant tag. **Tag All** All frames, whether classified to the Port VLAN or not, are transmitted with a tag. **Untag All** All frames, whether classified to the Port VLAN or not, are transmitted without a tag. This option is only available for ports in Hybrid mode.
Allowed VLANs	Ports in Trunk and Hybrid mode may control which VLANs they are allowed to become members of. Access ports can only be member of one VLAN, the Access VLAN. The field's syntax is identical to the syntax used in the Enabled VLANs field. By default, a Trunk or Hybrid port will become member of all VLANs, and is therefore set to **1-4095**. The field may be left empty, which means that the port will not become member of any VLANs.
Forbidden VLANs	A port may be configured to never become member of one or more VLANs. This is particularly useful when dynamic VLAN protocols like MVRP and GVRP must be prevented from dynamically adding ports to VLANs. The trick is to mark such VLANs as forbidden on the port in question. The syntax is identical to the syntax used in the Enabled VLANs field. By default, the field is left blank, which means that the port may become a member of all possible VLANs.

4.23.2 SVL

Shared VLAN Learning (SVL) allows for frames initially classified to a particular VLAN (based on Port VLAN ID or VLAN tag information) be bridged on a shared VLAN. In SVL, two or more VLANs are grouped to share common source address information in the MAC table. The common entry in the MAC table is identified by a Filter ID (FID). SVL is useful for configuration of more complex, asymmetrical cross-VLAN traffic patterns, like E-TREE (Rooted-Multipoint) and Multi-netted Server.

In SVL, one or more VLANs map to a Filter ID (FID). By default, there is a one-to-one mapping from VLAN to FID, in which case the switch acts as an IVL bridge, but with SVL multiple VLANs may share the same MAC address table entries.

The alternative VLAN learning mode is Independent VLAN Learning (IVL). The default VLAN learning mode is IVL and not all switches support SVL.

VLAN port settings and status can be viewed at Monitor→VLANs.

Shared VLAN Learning Configuration

Delete	FID	VLANs
Delete	1	

Parameter	Description
Delete	A previously allocated FID can be deleted by the use of this button.
FID	The Filter ID (FID) is the ID that VLANs get learned on in the MAC table when SVL is in effect. No two rows in the table can have the same FID and the FID must be a number **between 1 and 4095**.
VLANs	List of VLANs mapped into FID. The syntax is as follows: Individual VLANs are separated by commas. Ranges are specified with a

	dash separating the lower and upper bound. The following example will map VLANs 1, 10, 11, 12, 13, 200, and 300: 1,10-13,200,300. Spaces are allowed in between the delimiters. The range of valid VLANs is 1 to 4095. The same VLAN can only be a member of one FID. A message will be displayed if one VLAN is grouped into two or more FIDs. All VLANs must map to a particular FID, and by default VLAN x maps to FID x. This implies that if FID x is defined, then VLAN x is implicitly a member of FID x unless it is specified for another FID. If FID x doesn't exist, a confirmation message will be displayed, asking whether to continue adding VLAN x implicitly to FID x.

4.24 VLAN Translation

VLAN translation maps traffic from one VLAN ID to another as frames pass through the switch. It allows interconnection of networks using different VLAN schemes without changing end-device configurations, supporting flexible network integration and service separation.

4.24.1 Port to Group Configuration

This page allows you to configure switch Ports to use a given VLAN Translation Mapping Group. This will enable all VLAN Translation mappings of that group (if any) on the selected switch port.

VLAN Translation Port to Group Configuration

Interface	Group Configuration	
	Default	Group ID
*	☐	<>
Gi 1/1	☐	1
Gi 1/2	☐	2
Gi 1/3	☐	3
Gi 1/4	☐	4
Gi 1/5	☐	5
Gi 1/6	☐	6
Gi 1/7	☐	7
Gi 1/8	☐	8
10G 1/1	☐	9
10G 1/2	☐	10
10G 1/3	☐	11
10G 1/4	☐	12

Parameter	Description
Interface	The Interface column shows the list of port interfaces for which you can configure the VLAN Translation Mapping Group.
Default	To set the switch port to use the default VLAN Translation Group click the checkbox and press Save.
Group ID	The VLAN Translation mappings are organized into Groups, identified by the Group ID. In this way, a port is configured to use a number of VLAN Translation mappings easily by simply configuring it to use a given group. Then number of possible groups in a switch is equal to the number of ports present in this switch. A port can be configured to use any of the groups, but only one at any given time. Multiple ports can be configured to use the same group. *Note: By default, each port is set to use the group with Group ID equal to the port number. For example, port #3 is by default set to use group with Group ID = 3.*

4.24.2 VLAN Translation Mappings

This page allows you to view current or create new mappings of VLANs to Translated VLANs and organize these mappings into global Groups. Existing entries cannot be modified, because there is a risk that attempting to do so causes conflicts. If you need to modify an existing entry, add a new with the modified, then click the existing entry's "Delete" button followed by a click on "Save".

VLAN Translation Mapping Configuration

Delete	Group ID	Direction	VLAN ID	Translated VLAN ID
Delete	1	Both	1	1

Parameter	Description
Delete	Click this button to delete an entry in the mapping table. The action happens immediately and without the user being prompted. If the entry was added with a click on the Add button, it is simply removed again.
Group ID	The VLAN Translation mappings are organized into Groups, identified by the Group ID. In this way, a port is configured to use a number of VLAN Translation mappings easily by simply configuring it to use a given group. The number of possible groups in a switch is equal to the number of ports present in this switch. A port can be configured to use any of the groups, but only one at any given time. Multiple ports can be configured to use the same group. Note: By default, each port is set to use the group with Group ID equal to the port number. For example, port #3 is by default set to use group with Group ID = 3.
Direction	Indicates the direction of the VLAN Translation as seen from within the switch. The direction can be 'Ingress', where the translation takes place on the VLAN ID of frames entering the switch port, 'Egress', where the translation takes place on the VLAN ID of frames exiting the switch port, or 'Both', where the translation takes place on both of the above directions.
VLAN ID	Indicates the VLAN ID of the mapping (i.e. 'source' VLAN). A valid VLAN ID ranges from **`1 to 4095`**.
Translated VLAN ID	Indicates the translated VLAN ID to which a VLAN ID of a frame will be translated. A valid translated VLAN ID ranges from **`1 to 4095`**.

4.25 Private VLANs

In a private VLAN, PVLANs provide layer 2 isolation between ports within the same broadcast domain. Isolated ports configured as part of PVLAN cannot communicate with each other. Member ports of a PVLAN can communicate with each other.

Private VLANs are based on the source port mask, and there are no connections to VLANs. This means that VLAN IDs and Private VLAN IDs can be identical.

4.25.1 Membership

A port must be a member of both a VLAN and a Private VLAN to be able to forward packets. By default, all ports are VLAN unaware and members of VLAN 1 and Private VLAN 1.

A VLAN unaware port can only be a member of one VLAN, but it can be a member of multiple Private VLANs.

Private VLAN Membership Configuration

		Port Members											
Delete	PVLAN ID	1	2	3	4	5	6	7	8	9	10	11	12
☐	1	☑	☑	☑	☑	☑	☑	☑	☑	☑	☑	☑	☑

Parameter	Description
Delete	To delete a private VLAN entry, check this box. The entry will be deleted during the next save.

Private VLAN ID	Indicates the ID of this particular private VLAN. The allowed range for a private VLAN ID is the same as the switch port number range. Any values outside this range are not accepted, and a warning message appears.
Port Members	A row of check boxes for each port is displayed for each private VLAN ID. To include a port in a Private VLAN, check the box. To remove or exclude the port from the Private VLAN, make sure the box is unchecked. By default, no ports are members, and all boxes are unchecked.

4.25.2 Port Isolation

This page is used for enabling or disabling port isolation on ports in a Private VLAN.
Additional to isolating ports in different private VLAN, a port member of a VLAN can be isolated to other isolated ports on the same VLAN and Private VLAN by using Port Isolation.

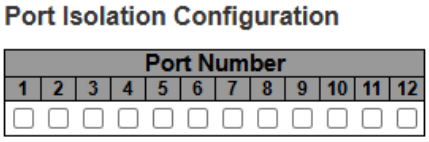

Parameter	Description
Port Members	A check box is provided for each port of a private VLAN. When checked, port isolation is enabled on that port. When unchecked, port isolation is disabled on that port. By default, port isolation is disabled on all ports.

4.26 VCL

A VLAN Control List (VCL) controls VLAN-related traffic and applies filtering rules by assigning VLAN IDs based on frame attributes like MAC addresses, protocols or IP-subnets. It is used to permit or deny VLAN traffic, providing additional control and security over VLAN forwarding behavior.

4.26.1 MAC-based VLAN

The MAC address to VLAN ID mappings can be configured here. This page allows adding and deleting MAC-based VLAN Classification List entries and assigning the entries to different ports. The maximum possible MAC to VLAN ID mapping entries are limited to 256.

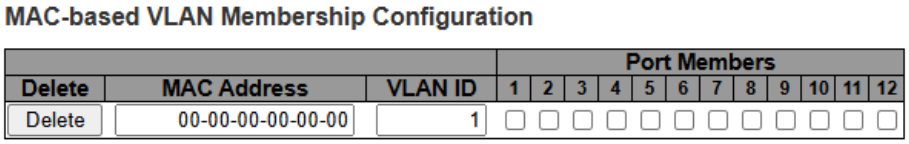

Parameter	Description
Delete	To delete a MAC to VLAN ID mapping entry, check this box and press save. The entry will be deleted in the stack.
MAC Address	Indicates the MAC address of the mapping. Any unicast MAC address can be used to configure the mapping. No broadcast or multicast MAC addresses are allowed.
VLAN ID	Indicates the VLAN ID the above MAC will be mapped to. Legal values for a VLAN ID are **1 to 4095.**
Port Members	A row of check boxes for each port is displayed for each MAC to VLAN ID mapping entry. To include a port in the mapping, check the box. To remove or exclude the port from the mapping, make sure the box is unchecked. By default, no ports are members, and all boxes are unchecked.

4.26.2 Protocol-based VLAN

A protocol-based VLAN assigns traffic to VLANs based on the Layer 3 protocol type. This allows the switch to segregate and manage traffic by protocol, improving security, organization, and traffic control within the network.

4.26.2.1 Protocol to Group

This page allows you to add new Protocol to Group Name (each protocol can be part of only one Group) mapping entries as well as allow you to see and delete already mapped entries for the switch. The maximum possible Protocol to Group mappings are limited to 128.

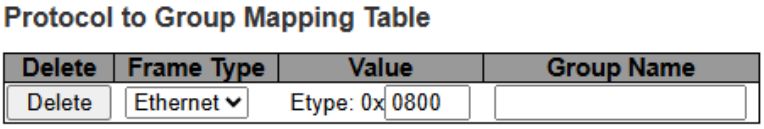

Parameter	Description
Delete	To delete a Protocol to Group Name map entry, check this box. The entry will be deleted from the switch during the next Save.
Frame Type	Frame Type can have one of the following values: **Ethernet** **LLC** **SNAP** ***Note**: When changing the Frame type field, the valid value of the following text field will vary depending on the new frame type you selected.*
Value	Valid value that can be entered in this text field depends on the option selected from the preceding Frame Type selection menu. Below are the criteria for the three different Frame Types: 1. **Ethernet:** Value in the text field when Ethernet is selected as a Frame Type is called etype. Valid values for etype range between 0x0600 and 0xffff 2. **LLC:** Valid value in this case is comprised of two different sub-values. a. **DSAP:** 1-byte long string (0x00-0xff) b. **SSAP:** 1-byte long string (0x00-0xff) 3. **SNAP:** Valid value in this case is also comprised of two different sub-values. a. **OUI:** OUI (Organizationally Unique Identifier) is a parameter in the format of xx-xx-xx where each pair (xx) in the string is a hexadecimal value ranging between 0x00 and 0xff. b. **PID:** PID (Protocol ID). If OUI is hexadecimal 000000, then the protocol ID is the Ethernet type (EtherType) field value for the protocol running on top of SNAP; if OUI is an OUI for a particular organization, the protocol ID is a value assigned by that organization to the protocol running on top of SNAP. In other words, if the value of OUI field is 00-00-00 then the value of PID will be etype (0x0600-0xffff) and if the value of OUI is other than 00-00-00 then valid values of PID will be any value between 0x0000 and 0xffff.
Group Name	A valid Group Name is a 16-character long string, unique for every entry, which consists of a combination of alphabets (a-z or A-Z) and integers (0-9). *Note: Special characters and underscores (_) are not allowed.*

4.26.2.2 Group to VLAN

This page allows you to map a Group Name (already configured or to be configured in the future) to a VLAN for the switch. The maximum possible Group to VLAN mappings are limited to 256.

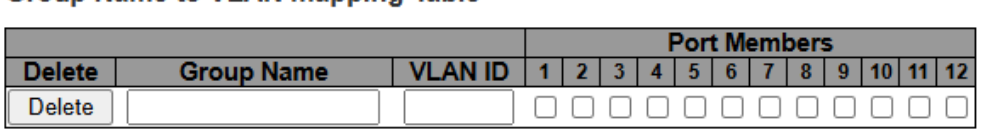

Parameter	Description
Delete	To delete a Group Name to VLAN mapping, check this box. The entry will be deleted from the switch during the next Save.
Group Name	A valid Group Name is a string, at the most 16 characters long, which consists of a combination of alphabets (a-z or A-Z) and integers (0-9) with no special characters allowed. You may either use a Group that already includes one or more protocols (see Protocol to Group mappings) or create a Group to VLAN ID mapping that will become active the moment you add one or more protocols inside that Group. Furthermore, the Group to VLAN ID mapping is not unique, as long as the port lists of these mappings are mutually exclusive (e.g. Group1 can be mapped to VID 1 on port#1 and to VID 2 on port#2).
VLAN ID	Indicates the VLAN ID to which the Group Name will be mapped. A valid VLAN ID ranges from **1 to 4095**.
Port Members	A row of check boxes for each port is displayed for each Group Name to VLAN ID mapping. To include a port in the mapping, check the box. To remove or exclude the port from the mapping, make sure the box is unchecked. By default, no ports are members, and all boxes are unchecked.

4.26.3 IP Subnet-based VLAN

The IP subnet to VLAN ID mappings can be configured here. This page allows adding, updating and deleting IP subnet to VLAN ID mapping entries and assigning them to different ports. The maximum possible IP subnet to VLAN ID mappings are limited to 128.

IP Subnet-based VLAN Membership Configuration

				Port Members											
Delete	IP Address	Mask Length	VLAN ID	1	2	3	4	5	6	7	8	9	10	11	12
Delete	0.0.0.0	24	1	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐

Parameter	Description
Delete	To delete a mapping, check this box and press save. The entry will be deleted in the stack.
IP Address	Indicates the subnet's IP address (Any of the subnet's host addresses can be also provided here, the application will convert it automatically).
Mask Length	Indicates the subnet's mask length.
VLAN ID	Indicates the VLAN ID the subnet will be mapped to. IP Subnet to VLAN ID is a unique matching. Legal values for the VLAN ID are **1 to 4095**.
Port Members	A row of check boxes for each port is displayed for each IP subnet to VLAN ID mapping entry. To include a port in a mapping, simply check the box. To remove or exclude the port from the mapping, make sure the box is unchecked. By default, no ports are members and all boxes are unchecked.

4.27 Voice VLAN

Voice VLAN is VLAN configured especially for voice traffic. By adding ports with voice devices attached to voice VLAN, we can perform QoS-related configuration for voice data, ensuring the transmission priority of voice traffic and voice quality.

4.27.1 Configuration

It is recommended that there be two VLANs on a port - one for voice, one for data. Before connecting the IP device to the switch, the IP phone should configure the voice VLAN ID correctly. It should be configured through its own GUI.

Voice VLAN Configuration

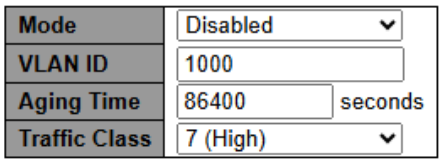

Parameter	Description
Mode	Indicates the Voice VLAN mode operation. We must disable MSTP feature before we enable Voice VLAN. It can avoid the conflict of ingress filtering. Possible modes are: **Enabled**: Enable Voice VLAN mode operation. **Disabled**: Disable Voice VLAN mode operation.
VLAN ID	Indicates the Voice VLAN ID. It should be a unique VLAN ID in the system and cannot equal each port PVID. It is a conflict in configuration if the value equals management VID, MVR VID, PVID etc. The allowed range is **1 to 4095**.
Aging Time	Indicates the Voice VLAN secure learning aging time. The allowed range is **10 to 10000000** seconds. It is used when security mode or auto detect mode is enabled. In other cases, it will be based on hardware aging time. The actual aging time will be situated between the **[age_time; 2 * age_time]** interval.

Port Configuration

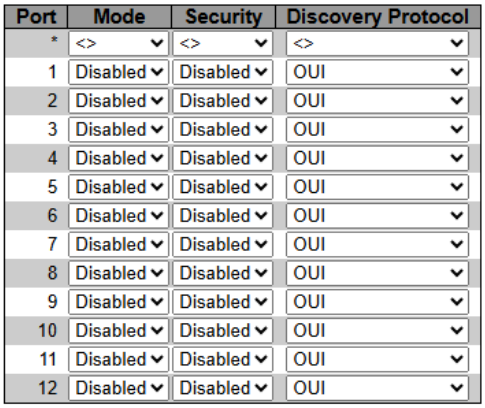

Parameter	Description
Port	The logical switch port to configure.
Mode	Indicates the Voice VLAN port mode. Possible port modes are: **Disabled**: Disjoin from Voice VLAN. **Auto**: Enable auto detect mode. It detects whether there is VoIP phone attached to the specific port and configures the Voice VLAN members automatically. **Forced**: Force joining the Voice VLAN.
Security	Indicates the Voice VLAN port security mode. When the function is enabled, all non-telephonic MAC addresses in the Voice VLAN will be blocked for 10 seconds. Possible port modes are: **Enabled**: Enable Voice VLAN security mode operation. **Disabled**: Disable Voice VLAN security mode operation.
Port Discovery Protocol	Indicates the Voice VLAN port discovery protocol. It will only work when auto detect mode is enabled. We should enable LLDP feature before configuring discovery protocol to "LLDP" or "Both". Changing the discovery protocol to "OUI" or "LLDP" will restart auto detect process. Possible discovery protocols are: **OUI**: Detect telephony device by OUI address. **LLDP**: Detect telephony device by LLDP. **Both**: Both OUI and LLDP.

4.27.2 OUI

OUI is an organizationally unique identifier. An OUI address is a globally unique identifier assigned to a vendor by IEEE. You can determine which vendor a device belongs to according to the OUI address which forms the first 24 bits of a MAC address.
The maximum number of entries is 16. Modifying the OUI table will restart auto detection of OUI process.

Voice VLAN OUI Table

Delete	Telephony OUI	Description
☐	00-01-e3	Siemens AG phones
☐	00-03-6b	Cisco phones
☐	00-0f-e2	H3C phones
☐	00-60-b9	Philips and NEC AG phones
☐	00-d0-1e	Pingtel phones
☐	00-e0-75	Polycom phones
☐	00-e0-bb	3Com phones

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
Telephony OUI	A telephony OUI address is a globally unique identifier assigned to a vendor by IEEE. It must be 6 characters long and the input format is "**xx-xx-xx**" (x is a hexadecimal digit).
Description	The description of OUI address. Normally, it describes which vendor telephony device it belongs to. The allowed string length is 0 to 32.

4.28 QoS

Quality of Service (QoS) is a method to guarantee a bandwidth relationship between individual applications or protocols.
A communications network transports a multitude of applications and data, including high-quality video and delay-sensitive data such as real-time voice. Networks must provide secure, predictable, measurable, and sometimes guaranteed services.

Achieving the required QoS becomes the key to a successful end-to-end business solution. Therefore, QoS provides set of techniques to manage network resources and ensure that critical data traffic will be handled with higher priority during periods of increased network traffic.
The QoS statistics can be viewed at Monitor→Ports→QoS Statistics.

4.28.1 Port Classification

This page allows you to configure the basic QoS Classification settings for all switch ports. Ingress traffic is assigned to a priority level per port for multiple different QoS classification methods.

QoS Port Classification

Port	Ingress								Egress
	CoS	DPL	PCP	DEI	CoS ID	Tag Class.	DSCP Based	Map	Map
*	<>	<>	<>	<>	<>				
1	0	0	0	0	0	Disabled			
2	0	0	0	0	0	Disabled			
3	0	0	0	0	0	Disabled			
4	0	0	0	0	0	Disabled			
5	0	0	0	0	0	Disabled			
6	0	0	0	0	0	Disabled			
7	0	0	0	0	0	Disabled			
8	0	0	0	0	0	Disabled			
9	0	0	0	0	0	Disabled			
10	0	0	0	0	0	Disabled			
11	0	0	0	0	0	Disabled			
12	0	0	0	0	0	Disabled			

Parameter	Description
Port	The port number for which the configuration below applies.
CoS	Controls the default Class of Service value. All frames are classified to a CoS. There is a one to one mapping between CoS, queue and priority. A CoS of 0 (zero) has the lowest priority. If the port is VLAN aware, the frame is tagged and Tag Class. is enabled, then the frame is classified to a CoS that is mapped from the PCP and DEI value in the tag. Otherwise, the frame is classified to the default CoS. The classified CoS can be overruled by a QCL entry. ***Note**: If the default CoS has been dynamically changed, then the actual default CoS is shown in parentheses after the configured default CoS.*
DPL	Controls the default DPL value (Drop Precedence Level). If the port is VLAN aware, the frame is tagged and Tag Class. is enabled, then the frame is classified to a DPL that is mapped from the PCP and DEI value in the tag. Otherwise, the frame is classified to the default DPL. The classified DPL can be overruled by a QCL entry.
PCP	Controls the default PCP value. If the port is VLAN aware and the frame is tagged, then the frame is classified to the PCP value in the tag. Otherwise, the frame is classified to the default PCP value.
DEI	Controls the default DEI value. If the port is VLAN aware and the frame is tagged, then the frame is classified to the DEI value in the tag. Otherwise, the frame is classified to the default DEI value.
CoS ID	Controls the default CoS ID value. Every incoming frame is classified to a CoS ID, which later can be used as basis for rewriting of different parts of the frame.
Tag Class.	Shows the classification mode for tagged frames on this port. **`Disabled`**: Use default CoS and DPL for tagged frames. **`Enabled`**: Use mapped versions of PCP and DEI for tagged frames. Click on the mode to configure the mode and/or mapping. ***Note**: This setting has no effect if the port is VLAN unaware. Tagged frames received on VLAN unaware ports are always classified to the default CoS and DPL.*
DSCP Based	Select to Enable DSCP Based QoS Ingress Port Classification.
Ingress Map	Controls the Ingress Map selection through the Map ID. The Ingress Map ID ranges from **`0 to 63`**. An empty field indicates no map selection. Ingress Map can be configured at Configuration→QoS→Ingress Map.
Egress Map	Controls the Egress Map selection through the Map ID. The Egress Map ID ranges from **`0 to 127`**. An empty field indicates no map selection. Egress Map can be configured at Configuration→QoS→Egress Map.

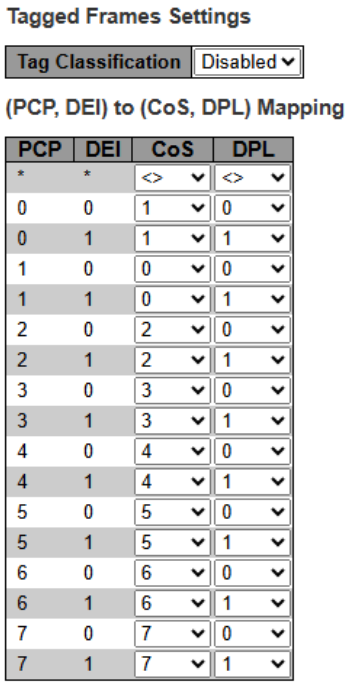

Parameter	Description
Tag Classification	Controls the classification mode for tagged frames on this port. **Disabled**: Use default CoS and DPL for tagged frames. **Enabled**: Use mapped versions of PCP and DEI for tagged frames.
(PCP, DEI) to (CoS, DPL) Mapping	Controls the mapping of the classified (PCP, DEI) to (CoS, DPL) values when Tag Classification is set to **Enabled**.

4.28.2 Port Policing

A policer can limit the bandwidth of received frames. It takes effect before received frames are added to the ingress queue. Port Policing can set the bandwidth for each port individually.

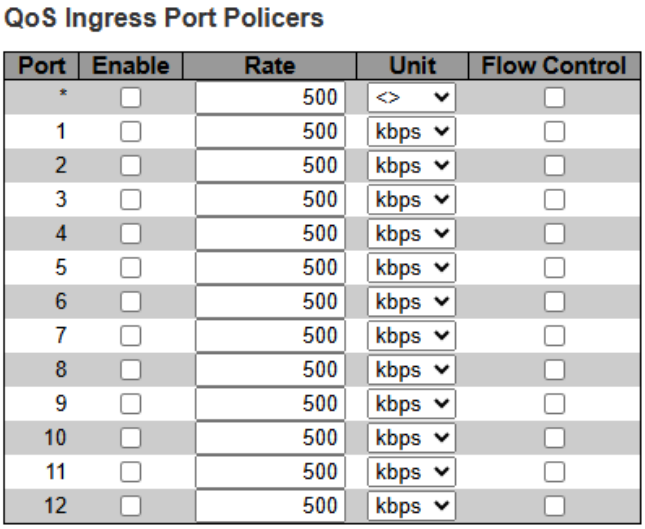

Parameter	Description
Port	The port number for which the configuration below applies.

Enable	Enable or disable the port policer for this switch port.
Rate	Controls the rate for the port policer. This value is restricted to **10-13128147** when *Unit* is **kbps or fps**, and **1-13128** when *Unit* is **Mbps or kfps**. The rate is internally rounded up to the nearest value supported by the port policer.
Unit	Controls the unit of measure for the port policer rate as kbps, Mbps, fps or kfps.
Flow Control	If flow control is enabled and the port is in flow control mode, then pause frames are sent instead of discarding frames.

4.28.3 Queue Policing

Queue Policing can set the bandwidth limit for each Queue individually.
The QoS Queuing Counter can be viewed at Monitor→Ports→QoS Statistics.

QoS Ingress Queue Policers

Port	Queue 0			Queue 1	Queue 2	Queue 3	Queue 4	Queue 5	Queue 6	Queue 7
	E	Rate	Unit	Enable	Enable	Enable	Enable	Enable	Enable	Enable
*	☑	500	<>	☐	☐	☐	☐	☐	☐	☐
1	☑	500	kbps	☐	☐	☐	☐	☐	☐	☐
2	☑	500	kbps	☐	☐	☐	☐	☐	☐	☐
3	☑	500	kbps	☐	☐	☐	☐	☐	☐	☐
4	☑	500	kbps	☐	☐	☐	☐	☐	☐	☐
5	☑	500	kbps	☐	☐	☐	☐	☐	☐	☐
6	☑	500	kbps	☐	☐	☐	☐	☐	☐	☐
7	☑	500	kbps	☐	☐	☐	☐	☐	☐	☐
8	☑	500	kbps	☐	☐	☐	☐	☐	☐	☐
9	☑	500	kbps	☐	☐	☐	☐	☐	☐	☐
10	☑	500	kbps	☐	☐	☐	☐	☐	☐	☐
11	☑	500	kbps	☐	☐	☐	☐	☐	☐	☐
12	☑	500	kbps	☐	☐	☐	☐	☐	☐	☐

Parameter	Description
Port	The port number for which the configuration below applies.
Enable / E	Enable or disable the queue policer for this switch port.
Rate	Controls the rate for the port policer. This value is restricted to **25-13128147** when *Unit* is **kbps**, and **1-13128** when *Unit* is **Mbps**. The rate is internally rounded up to the nearest value supported by the port policer.
Unit	Controls the unit of measure for the queue policer rate as kbps or Mbps. This field is only shown if at least one of the queue policers are enabled.

4.28.4 Port Scheduler and Shaper

The two pages *Port Scheduler* and *Port Shaping* provide an overview of the QoS Egress Port Scheduler and Port Shaper settings for all switch ports. When clicking a port ID, they both lead to the same respective configuration page for that port.

QoS Egress Port Schedulers

Port	Mode	Weight							
		Q0	Q1	Q2	Q3	Q4	Q5	Q6	Q7
1	Strict Priority	-	-	-	-	-	-	-	-
2	Strict Priority	-	-	-	-	-	-	-	-
3	Strict Priority	-	-	-	-	-	-	-	-
4	Strict Priority	-	-	-	-	-	-	-	-
5	Strict Priority	-	-	-	-	-	-	-	-
6	Strict Priority	-	-	-	-	-	-	-	-
7	Strict Priority	-	-	-	-	-	-	-	-
8	Strict Priority	-	-	-	-	-	-	-	-
9	Strict Priority	-	-	-	-	-	-	-	-
10	Strict Priority	-	-	-	-	-	-	-	-
11	Strict Priority	-	-	-	-	-	-	-	-
12	Strict Priority	-	-	-	-	-	-	-	-

QoS Egress Port Shapers

Port	Shapers								
	Q0	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Port
1	-	-	-	-	-	-	-	-	-
2	-	-	-	-	-	-	-	-	-
3	-	-	-	-	-	-	-	-	-
4	-	-	-	-	-	-	-	-	-
5	-	-	-	-	-	-	-	-	-
6	-	-	-	-	-	-	-	-	-
7	-	-	-	-	-	-	-	-	-
8	-	-	-	-	-	-	-	-	-
9	-	-	-	-	-	-	-	-	-
10	-	-	-	-	-	-	-	-	-
11	-	-	-	-	-	-	-	-	-
12	-	-	-	-	-	-	-	-	-

Parameter	Description
Port	The logical port for the settings contained in the same row. Click on the port number in order to configure the schedulers.
Mode	Shows the scheduling mode for this port.
Weight Qn	Shows the weight for this queue and port.
Shaper Qn	Shows "**-**" for disabled or actual queue shaper rate - **`e.g. "800 Mbps"`**.
Port	Shows "**-**" for disabled or actual port shaper rate - **`e.g. "800 Mbps"`**.

QoS Egress Port Scheduler and Shapers Port 1

Scheduler Mode: 4 Queues Weighted

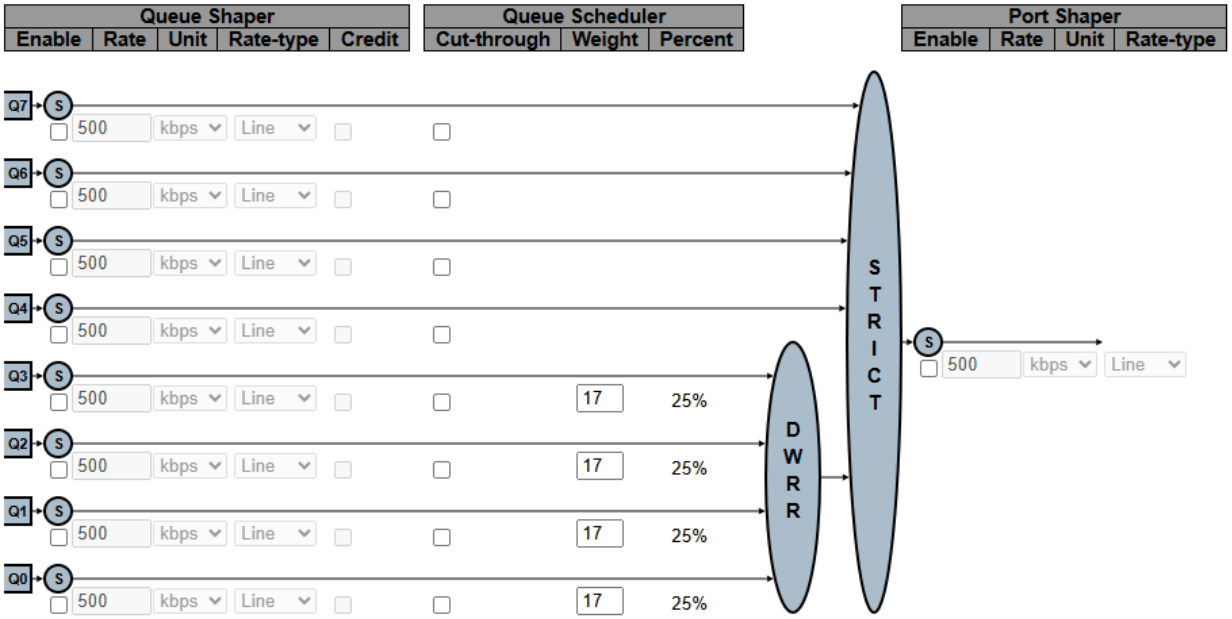

Parameter	Description
Scheduler Mode	The logical port for the settings contained in the same row. Click on the port number in order to configure the schedulers.
Queue Shaper Enable	Controls whether the queue shaper is enabled for this queue on this switch port.
Queue Shaper Rate	Controls the rate for the queue shaper. This value is restricted to **100-13107100** when *Unit* is **kbps**, and **1-13107** when *Unit* is **Mbps**. The rate is internally rounded up to the nearest value supported by the queue shaper.
Queue Shaper Unit	Controls the unit of measure for the queue shaper rate as **kbps or Mbps**.
Queue Shaper Rate-type	The rate type of the queue shaper. The allowed values are: **Line**: Specify that this shaper operates on line rate. **Data**: Specify that this shaper operates on data rate.
Queue Shaper Credit	Controls whether the queue has credit-based shaper enabled.
Queue Scheduler Cut-through	Controls whether the queue has cut-through enabled.
Queue Scheduler Preemption	Controls whether the queue has frame preemption enabled.
Queue Scheduler Weight	Controls the weight for this queue. This value is restricted to **1-100**. This parameter is only shown if *Scheduler Mode* is set to **Weighted**.
Queue Scheduler Percent	Shows the weight in percent for this queue. This parameter is only shown if *Scheduler Mode* is set to **Weighted**.
Port Shaper Enable	Controls whether the port shaper is enabled for this switch port.
Port Shaper Rate	Controls the rate for the port shaper. This value is restricted to **100-13107100** when *Unit* is **kbps**, and **1-13107** when *Unit* is **Mbps**. The rate is internally rounded up to the nearest value supported by the port shaper.
Port Shaper Unit	Controls the unit of measure for the port shaper rate as **kbps or Mbps**.
Port Shaper Rate-type	The rate type of the port shaper. The allowed values are: **Line**: Specify that this shaper operates on line rate. **Data**: Specify that this shaper operates on data rate.

4.28.5 Port Tag Remarking

QoS egress port tag remarking modifies the priority or QoS marking of VLAN tags on outgoing frames. It allows the switch to adjust traffic priority based on QoS policies before forwarding packets to downstream devices. This page provides an overview of QoS Egress Port Tag Remarking for all switch ports. It can be configured for a specific port by clicking on the port number.

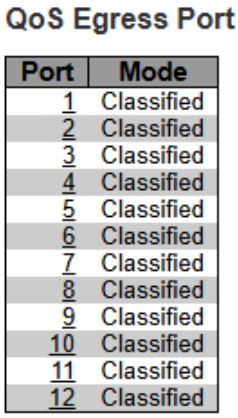

QoS Egress Port Tag Remarking

Port	Mode
1	Classified
2	Classified
3	Classified
4	Classified
5	Classified
6	Classified
7	Classified
8	Classified
9	Classified
10	Classified
11	Classified
12	Classified

Parameter	Description
Port	The logical port for the settings contained in the same row. Click on the port number in order to configure tag remarking.
Mode	Shows the tag remarking mode for this port. **Classified**: Use classified PCP/DEI values. **Default**: Use default PCP/DEI values. **Mapped**: Use mapped versions of CoS and DPL.

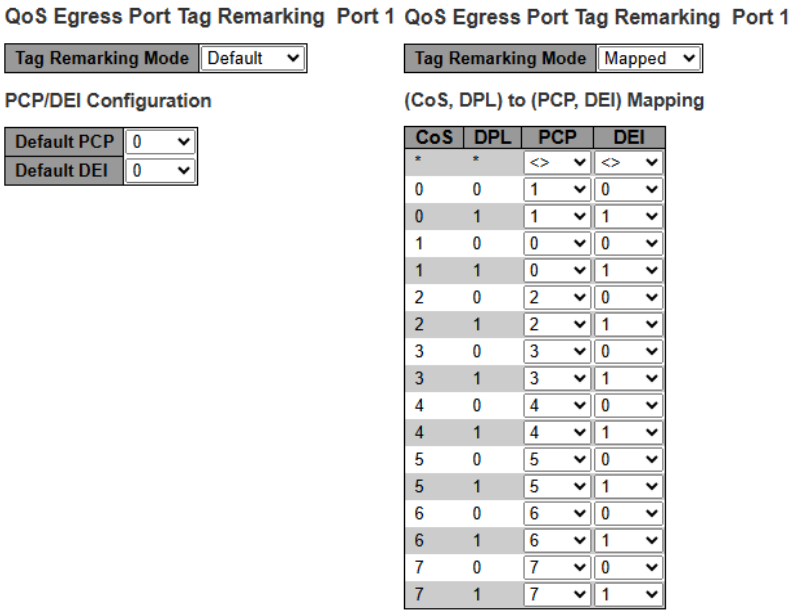

QoS Egress Port Tag Remarking Port 1

Tag Remarking Mode: Default

PCP/DEI Configuration

Default PCP	0
Default DEI	0

QoS Egress Port Tag Remarking Port 1

Tag Remarking Mode: Mapped

(CoS, DPL) to (PCP, DEI) Mapping

CoS	DPL	PCP	DEI
*	*	<>	<>
0	0	1	0
0	1	1	1
1	0	0	0
1	1	0	1
2	0	2	0
2	1	2	1
3	0	3	0
3	1	3	1
4	0	4	0
4	1	4	1
5	0	5	0
5	1	5	1
6	0	6	0
6	1	6	1
7	0	7	0
7	1	7	1

Parameter	Description
Mode	Controls the tag remarking mode for this port. **Classified**: Use classified PCP/DEI values. **Default**: Use default PCP/DEI values. **Mapped**: Use mapped versions of CoS and DPL.

PCP/DEI Configuration	Controls the default PCP and DEI values used when the mode is set to **Default**.
(CoS, DPL) to (PCP, DEI) Mapping	Controls the mapping of the classified (CoS, DPL) to (PCP, DEI) values when the mode is set to **Mapped**.

4.28.6 DSCP

DSCP is a field in the IP header used to classify and prioritize network traffic for Quality of Service (QoS). It allows devices to apply differentiated handling, such as priority queuing or bandwidth allocation, based on the assigned DSCP value.
This page allows you to configure the basic QoS Port DSCP Configuration settings for all switch ports.

QoS Port DSCP Configuration

Port	Ingress		Egress
	Translate	Classify	Rewrite
*	☐	<>	<>
1	☐	Disable	Disable
2	☐	Disable	Disable
3	☐	Disable	Disable
4	☐	Disable	Disable
5	☐	Disable	Disable
6	☐	Disable	Disable
7	☐	Disable	Disable
8	☐	Disable	Disable
9	☐	Disable	Disable
10	☐	Disable	Disable
11	☐	Disable	Disable
12	☐	Disable	Disable

Parameter	Description
Port	The Port column shows the list of ports for which you can configure DSCP ingress and egress settings.
Ingress	In Ingress settings you can change ingress translation and classification settings for individual ports. There are two configuration parameters available in Ingress: **1. Translate** To Enable the Ingress Translation click the checkbox. **2. Classify** Classification for a port have 4 different values. **Disable**: No Ingress DSCP Classification. **DSCP=0**: Classify if incoming (or translated if enabled) DSCP is 0. **Selected**: Classify only selected DSCP for which classification is enabled as specified in DSCP Translation window for the specific DSCP. **All**: Classify all DSCP.
Egress	Port Egress Rewriting can be one of **Disable**: No Egress rewrite. **Enable**: Rewrite enabled without remapping. **Remap**: DSCP from analyzer is remapped and frame is remarked with remapped DSCP value.

4.28.7 DSCP-Based QoS

DSCP-based ingress classification assigns traffic priority based on the DSCP value in incoming IP packets. Each DSCP value can be mapped to an internal Class of Service (CoS) and Drop Precedence Level (DPL), which determine packet priority and congestion handling.

DSCP-Based QoS Ingress Classification

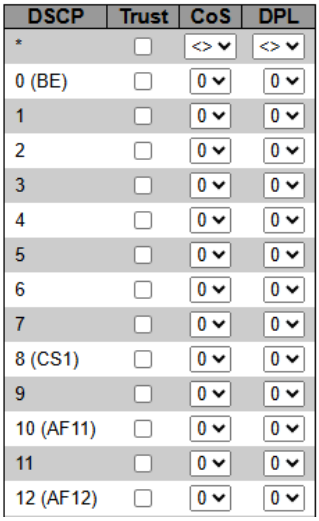

DSCP	Trust	CoS	DPL
*	☐	<>	<>
0 (BE)	☐	0	0
1	☐	0	0
2	☐	0	0
3	☐	0	0
4	☐	0	0
5	☐	0	0
6	☐	0	0
7	☐	0	0
8 (CS1)	☐	0	0
9	☐	0	0
10 (AF11)	☐	0	0
11	☐	0	0
12 (AF12)	☐	0	0

Parameter	Description
DSCP	supported DSCP values are **0 to 63**.
Trust	Controls whether a specific DSCP value is trusted. Only frames with trusted DSCP values are mapped to a specific CoS and DPL. Frames with untrusted DSCP values are treated as a non-IP frame.
CoS	CoS value can be any of **0 to 7**
DPL	Drop Precedence Level (**0-3**)

4.28.8 DSCP Translation

This page allows you to configure the basic QoS DSCP Translation settings for all switches. DSCP translation can be done in Ingress or Egress.

DSCP Translation

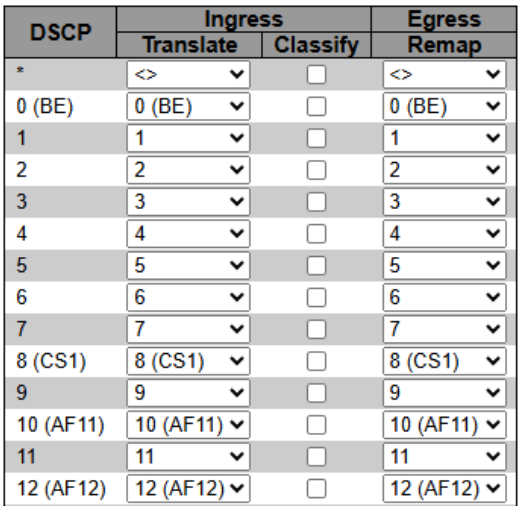

DSCP	Ingress		Egress
	Translate	Classify	Remap
*	<>	☐	<>
0 (BE)	0 (BE)	☐	0 (BE)
1	1	☐	1
2	2	☐	2
3	3	☐	3
4	4	☐	4
5	5	☐	5
6	6	☐	6
7	7	☐	7
8 (CS1)	8 (CS1)	☐	8 (CS1)
9	9	☐	9
10 (AF11)	10 (AF11)	☐	10 (AF11)
11	11	☐	11
12 (AF12)	12 (AF12)	☐	12 (AF12)

Parameter	Description
DSCP	Maximum number of supported DSCP values are 64 and valid DSCP value ranges from **0 to 63**.
Ingress	Ingress side DSCP can be first translated to new DSCP before using the DSCP for CoS and DPL map. There are two configuration parameters for DSCP Translation

	1. Translate DSCP at Ingress side can be translated to any of (0-63) DSCP values. **2. Classify** Click to enable Classification at Ingress side.
Egress Remap	Select the DSCP value from select menu to which you want to remap. DSCP value ranges from **0 to 63**.

4.28.9 DSCP Classification

This page allows you to configure the mapping of CoS and DPL to DSCP value.

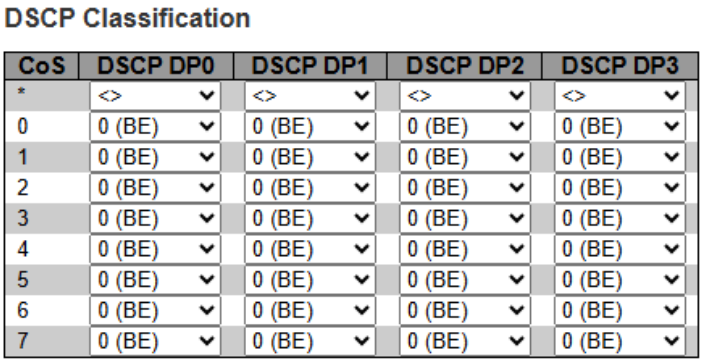

Parameter	Description
CoS	Maximum number of supported DSCP values are 64 and valid DSCP value ranges from **0 to 63**.
DSCP DP0	Select the classified DSCP value (**0-63**) for Drop Precedence Level 0.
DSCP DP1	Select the classified DSCP value (**0-63**) for Drop Precedence Level 1.
DSCP DP2	Select the classified DSCP value (**0-63**) for Drop Precedence Level 2.
DSCP DP3	Select the classified DSCP value (**0-63**) for Drop Precedence Level 3.

4.28.10 Ingress Map

This page shows a table of QoS Ingress Maps which is made up of individual map entries. Each entry has a key and an action. The key indicates which fields of the frame will be mapped to the fields specified by and according to the action. Each Map can hold a number of map rules, or mappings between possible keys and actions. Which of those rules will be applied depends on the selection of (Key-Type, Action-Type). Each row describes a user-defined map. The maximum number of Ingress Maps is **64**. Each Ingress Map uses a number of key-entries in a internal key mapping table which have **1004** key-entries available for configuration. The consumption of key-entries by Key Type are listed as table width in the Key-Type table below. A new Ingress Map can only be defined when there are sufficient free key-entries.
NOTE: The left table is just an overview of the configured maps. The user can add new ones or edit existing maps using the Add/Edit buttons, which will open the interface on the right. Click on the lowest plus sign (empty map entry) to add a new Ingress Map to the table.

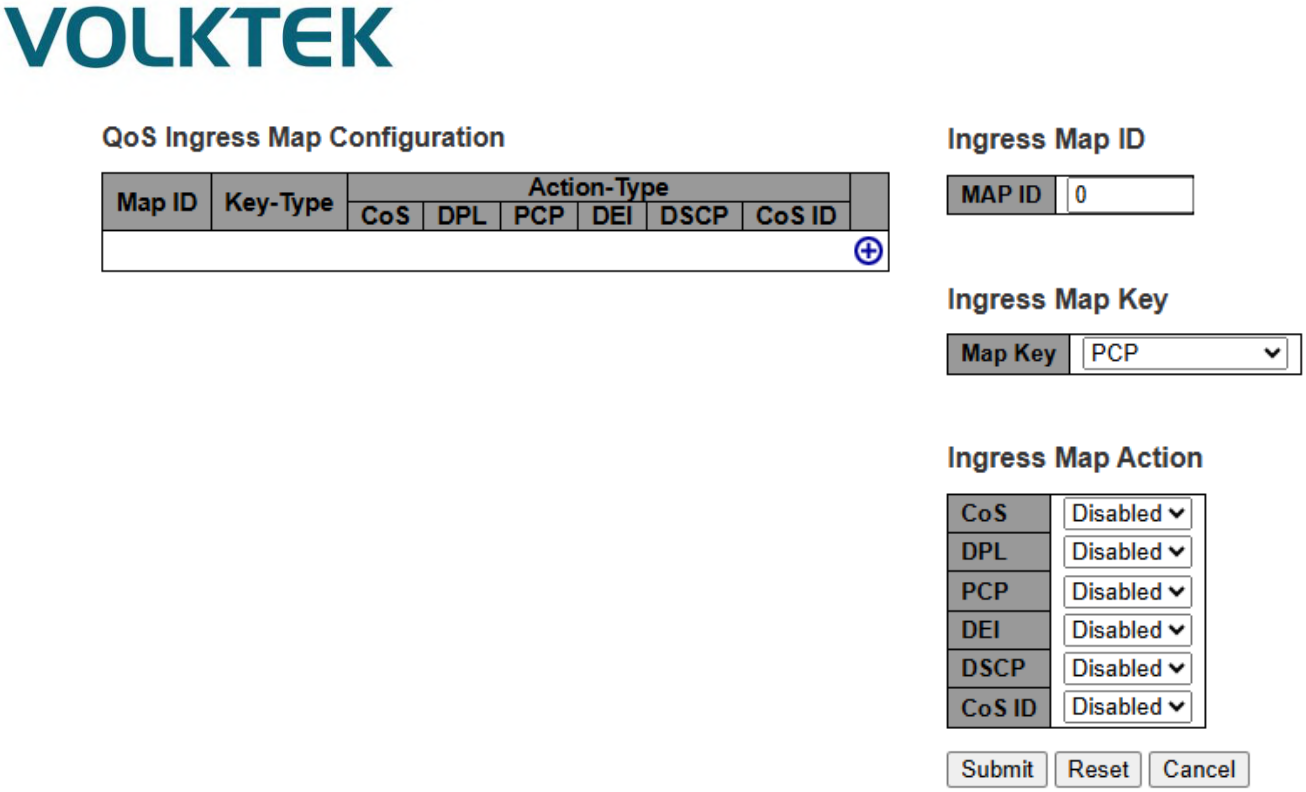

Parameter	Description
Map ID	Indicates the Map (unique) ID. Range is **0 to 63**. When in edit mode, this is non-configurable. However, it is possible to overwrite an existing mapping through the create mode.
Map Key	Indicates the Key type that will be used to filter the map rules when applying the map. As mentioned above, map rules can have various keys and this is to make a select set of them. Possible Key types are: **PCP**: Use PCP as key for tagged frames and none for the rest. **PCP - DEI**: Use PCP/DEI as key for tagged frames and none for the rest. **DSCP**: Use DSCP as key for IP frames and none for the rest. **DSCP - PCP - DEI**: Use DSCP as key for IP frames, PCP/DEI for tagged frames and none for the rest.
Map Action	Indicates the Action type that will be used to filter the map rules when applying the map. As mentioned above, map rules can have various actions available and this is to make a select set of them. Possible Action types are: **CoS**: Class of Service. **DPL**: Drop Precedence Level. **PCP**: Priority Code Point. **DEI**: Drop Eligible Indicator. **DSCP**: Differentiated Services Code Point. **CoS ID**: CoS ID.

4.28.11 Egress Map

This page shows a table of QoS Egress Maps which is made up of individual map entries. Each entry has a key and an action. The key indicates which fields of the frame will be mapped to the fields specified by and according to the action. Each Map can hold a number of map rules, or mappings between possible keys and actions. Which of those rules will be applied depends on the selection of (Key-Type, Action-Type). Each row describes a user-defined map. The maximum number of Egress Maps is **128**. Each Egress Map uses a number of key-entries in an internal key mapping table which have **960** key-entries available. The consumption of key-entries by Key Type are listed as table width in the Key-Type table below. A new Egress Map can only be defined when there are sufficient free key-entries.
NOTE: The left table is just an overview of the configured maps. The user can add new ones or edit existing

maps using the Add/Edit buttons, which will open the interface on the right. Click on the lowest plus sign (empty map entry) to add a new Ingress Map to the table.

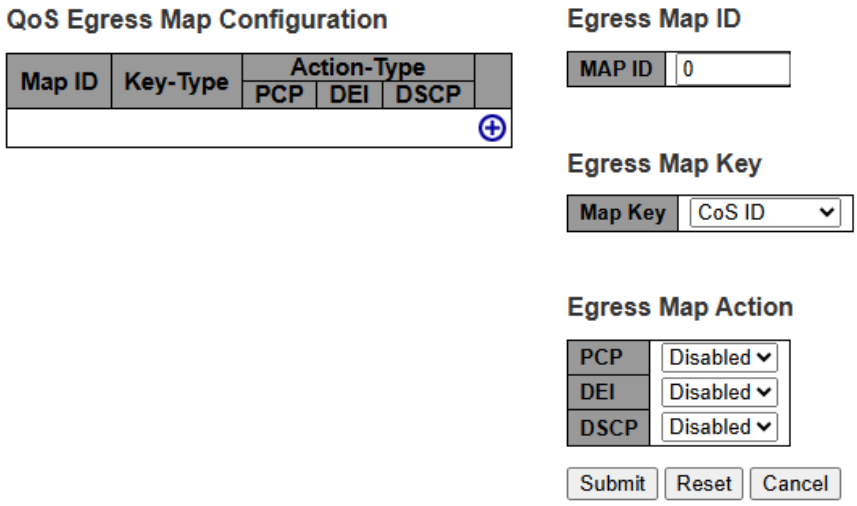

Parameter	Description
Map ID	Indicates the Map (unique) ID. Range is **0 to 127**. When in edit mode, this is non-configurable. However, it is possible to overwrite an existing mapping through the create mode.
Key Type / Map Key	Indicates the Key type that will be used to filter the map rules when applying the map. As mentioned above, map rules can have various keys and this is to make a select set of them. Possible Key types are: **CoS ID**: Use classified COS ID as key. **CoS ID - DPL**: Use classified COS ID and DPL as key. **DSCP**: Use classified DSCP as key. **DSCP - DPL:** Use classified DSCP and DPL as key.
Map Action	Indicates the Action type that will be used to filter the map rules when applying the map. As mentioned above, map rules can have various actions available, and this is to make a select set of them. Possible Action types are: **PCP**: Priority Code Point. **DEI**: Drop Eligible Indicator. **DSCP**: Differentiated Services Code Point.

4.28.12 QoS Control List

The QoS Control List (QCL) is a list of QoS Control Entries (QCEs).
A QCE is a combination of keys and actions. The keys can be configured to match specific parts of a frame and the actions can be configured to override the default classified values of e.g. CoS.
Each frame is compared against the QCEs in the list. The comparison starts with the first entry in the list and continues until there is a match between the frame and the key parameters or the end of the list is reached.
If there is a match between the frame and the keys, the frame will be reclassified according to the action parameters.
The maximum number of QCEs is 256. Click on the lowest plus sign to add a new QCE to the list.
The Control List status is displayed at Monitor→Ports→QCL Status.

QCE Configuration

Port Members											
1	2	3	4	5	6	7	8	9	10	11	12
☑	☑	☑	☑	☑	☑	☑	☑	☑	☑	☑	☑

Key Parameters

DMAC	Any			
SMAC	Specific	00-00-00-00-00-00		
Tag	Any			
VID	Range	From:	To	
PCP	Any			
DEI	Any			
Inner Tag	Any			
Inner VID	Range	From:	To	
Inner PCP	Any			
Inner DEI	Any			
Frame Type	Any			

Action Parameters

CoS	0
DPL	Default
DSCP	Default
PCP	Default
DEI	Default
Policy	
Ingress Map ID	

Parameter	Description
QCE	Indicates the QCE ID.
Port Members	Check the checkbox button to include the port in the QCL entry. By default all ports are included.
DMAC	Destination MAC address: Possible values are **'Unicast', 'Multicast', 'Broadcast', 'Specific' (xx-xx-xx-xx-xx-xx) or 'Any'**.
SMAC	Source MAC address**: xx-xx-xx-xx-xx-xx or 'Any'**.
Tag	Value of Tag field can be **'Untagged', 'Tagged', 'C-Tagged', 'S-Tagged' or 'Any'**.
VID	Valid value of VLAN ID can be any value in the range **1-4095 or 'Any'**; user can enter either a specific value or a range of VIDs.
PCP	Valid value PCP are **specific** (0, 1, 2, 3, 4, 5, 6, 7) or **range** (0-1, 2-3, 4-5, 6-7, 0-3, 4-7) or **'Any'**.
DEI	Valid value of DEI can be **'0', '1' or 'Any'**.
Inner Tag	Value of Inner Tag field can be **'Untagged', 'Tagged', 'C-Tagged', 'S-Tagged' or 'Any'**.
Inner VID	Valid value of Inner VLAN ID can be any value in the range **1-4095 or 'Any'**; user can enter either a specific value or a range of VIDs.
Inner PCP	Valid value of Inner PCP are **specific** (0, 1, 2, 3, 4, 5, 6, 7) or **range** (0-1, 2-3, 4-5, 6-7, 0-3, 4-7) or **'Any'**.
Inner DEI	Valid value of Inner DEI can be **'0', '1' or 'Any'**.
Frame Type	Frame Type can have any of the following values: **1. Any** Allow all types of frames. **2. EtherType** Valid Ether Type can be **0x600-0xFFFF** excluding 0x800(IPv4) and 0x86DD(IPv6) or **'Any'**. **3. LLC** **DSAP Address** Valid DSAP(Destination Service Access Point) can vary from 0x00 to 0xFF or 'Any'. **SSAP Address** Valid SSAP(Source Service Access Point) can vary from 0x00 to 0xFF or 'Any'. **Control** Valid Control field can vary from 0x00 to 0xFF or 'Any'. **4. SNAP** **PID** Valid PID (a.k.a Ether Type) can be 0x0000-0xFFFF or 'Any'.

	5. IPv4 **Protocol IP** protocol number: (0-255, 'TCP' or 'UDP') or 'Any'. **Source IP** Specific Source IP address in value/mask format or 'Any'. IP and Mask are in the format x.y.z.w where x, y, z, and w are decimal numbers between 0 and 255. When Mask is converted to a 32-bit binary string and read from left to right, all bits following the first zero must also be zero. **Destination IP** Specific Destination IP address in value/mask format or 'Any'. **IP Fragment** IPv4 frame fragmented option: 'Yes', 'No' or 'Any'. **DSCP** Diffserv Code Point value (DSCP): It can be a specific value, range of values or 'Any'. DSCP values are in the range 0-63 including BE, CS1-CS7, EF or AF11-AF43. **Sport** Source TCP/UDP port:(0-65535) or 'Any', specific or port range applicable for IP protocol UDP/TCP. **Dport** Destination TCP/UDP port:(0-65535) or 'Any', specific or port range applicable for IP protocol UDP/TCP. **6. IPv6** **Protocol** IP protocol number: (0-255, 'TCP' or 'UDP') or 'Any'. **Source IP** 32 LS bits of IPv6 source address in value/mask format or 'Any'. **Destination IP** Specific Destination IP address in value/mask format or 'Any'. **DSCP** Diffserv Code Point value (DSCP): It can be a specific value, range of values or 'Any'. DSCP values are in the range 0-63 including BE, CS1-CS7, EF or AF11-AF43. **Sport** Source TCP/UDP port:(0-65535) or 'Any', specific or port range applicable for IP protocol UDP/TCP. **Dport** Destination TCP/UDP port:(0-65535) or 'Any', specific or port range applicable for IP protocol UDP/TCP.
Action Parameters	**CoS** Class of Service: (0-7) or 'Default'. **DPL** Drop Precedence Level: (0-3) or 'Default'. **DSCP** DSCP: (0-63, BE, CS1-CS7, EF or AF11-AF43) or 'Default'. **PCP** PCP: (0-7) or 'Default'. Note: PCP and DEI cannot be set individually. **DEI** DEI: (0-1) or 'Default'. **Policy** ACL Policy number: (0-255) or 'Default' (empty field). **Ingress Map ID** Ingress Map ID: (0-63) or no Ingress Map (empty field). 'Default' means that the default classified value is not modified by this QCE.

4.28.13 Storm Policing

Global storm policers for the switch are configured in the table below.

The unicast storm policer, multicast storm policer, and a broadcast storm policer only affect flooded frames, i.e. frames with a (VLAN ID, DMAC) pair not present in the MAC Address table.

Global Storm Policer Configuration

Frame Type	Enable	Rate	Unit
Unicast	☐	10	fps
Multicast	☐	10	fps
Broadcast	☐	10	fps

Parameter	Description
Frame Type	The frame type for which the configuration below applies.
Enable	Enable or disable the global storm policer for the given frame type.
Rate	Controls the rate for the global storm policer. This value is restricted to **10-13128147** when *Unit* is **fps or kbps**, and **1-13128** when *Unit* is **kfps or Mbps**. The rate is internally rounded up to the nearest value supported by the global storm policer. Supported rates are divisible by 10 fps or 25 kbps.
Unit	Controls the unit of measure for the global storm policer rate as **fps, kfps, kbps or Mbps**.

Port storm policers for all switch ports are configured in the table below.

There are storm policers for known and unknown unicast frames, known and unknown broadcast frames and unknown (flooded) unicast, multicast and broadcast frames.

Port Storm Policer Configuration

Port	Unicast Frames			Broadcast Frames			Unknown Frames		
	Enable	Rate	Unit	Enable	Rate	Unit	Enable	Rate	Unit
*	☐	500	<>	☐	500	<>	☐	500	<>
1	☐	500	kbps	☐	500	kbps	☐	500	kbps
2	☐	500	kbps	☐	500	kbps	☐	500	kbps
3	☐	500	kbps	☐	500	kbps	☐	500	kbps
4	☐	500	kbps	☐	500	kbps	☐	500	kbps
5	☐	500	kbps	☐	500	kbps	☐	500	kbps
6	☐	500	kbps	☐	500	kbps	☐	500	kbps
7	☐	500	kbps	☐	500	kbps	☐	500	kbps
8	☐	500	kbps	☐	500	kbps	☐	500	kbps
9	☐	500	kbps	☐	500	kbps	☐	500	kbps
10	☐	500	kbps	☐	500	kbps	☐	500	kbps
11	☐	500	kbps	☐	500	kbps	☐	500	kbps
12	☐	500	kbps	☐	500	kbps	☐	500	kbps

Parameter	Description
Port	The port number for which the configuration below applies.
Enable	Enable or disable the storm policer for this switch port.
Rate	Controls the rate for the port storm policer. This value is restricted to **10-13128147** when *Unit* is **fps or kbps**, and **1-13128** when *Unit* is **kfps or Mbps**. The rate is internally rounded up to the nearest value supported by the port storm policer. Supported rates are divisible by 10 fps or 25 kbps.
Unit	Controls the unit of measure for the port storm policer rate as **fps, kfps, kbps or Mbps**.

4.29 TSN

Time-Sensitive Networking (TSN) is a set of IEEE 802.1 standards that extends Ethernet to support deterministic, low-latency, and highly reliable data transmission. TSN enables time-critical applications, such as industrial automation, automotive networks, and professional audio/video systems, to share the same Ethernet infrastructure with standard data traffic.

TSN achieves deterministic behavior through features such as time synchronization, traffic scheduling, resource reservation, and traffic shaping. These mechanisms ensure that critical traffic is delivered within defined latency and jitter requirements while maintaining high network utilization and interoperability with existing Ethernet technologies.

4.29.1 Streams

This page gives an overview of the configured streams, and it allows for deleting and editing existing streams as well as adding new.
TSN status information can be viewed at Monitor→TSN.

Stream Configuration Overview

Stream ID	DMAC	SMAC	Outer VLAN Tag	Inner VLAN Tag	Protocol	Attached Clients	Stream Collection	Warnings	
									⊕

Parameter	Description
Stream ID	The ID of the stream. The lower the ID, the higher precedence when matching. For example, if Stream #1 matches all frames with a multicast DMAC and Stream #2 matches a particular multicast DMAC, Stream #2 will never be hit.
DMAC	Indicates which types of destination MAC addresses are matched.
SMAC	Indicates whether a particular source MAC address (with an optional mask) or any source MAC address is matched.
Outer VLAN Tag	Indicates the required presence of an outer VLAN tag.
Inner VLAN Tag	Indicates the required presence of an inner VLAN tag.
Protocol	Indicates layer 3 and 4 matching options.
Attached Clients	Streams don't do anything by themselves. They are utilized by PSFP and FRER for matching particular flows. This field indicates which of those two protocols - if any - are utilizing a particular stream. The number in parenthesis shows the PSFP or FRER instance the stream is used on. A stream may be part of a stream collection, in which case no clients can be attached directly to the stream itself. They must attach to the stream collection. In this case, this fields shows a dash ("-"). See also next help topic. If no clients are attached to a stream, the stream will be installed into hardware anyway. This allows the user to create two streams, where the first e.g. matches a particular range of MAC address and the second matches the same range of MAC addresses plus some more. The second stream can them be used in FRER or PSFP to match all the MAC addresses of the second stream, except that from the first.
Stream Collection	If a stream is part of a stream collection, this shows the ID of the stream collection. The ID is a hyperlink to the stream collection configuration page. If the stream is not part of a stream collection, the field just shows "None".
Warnings	Configuration of a stream may result in configurational warnings. If it is yellow, at least one configurational warning is detected. Hover the mouse over the image to see a list of the warnings. If it is green, no configurational warnings are detected. For a description of the possible configurational warnings, refer to Stream Configuration Help.

Stream Configuration

This page appears when clicking the ⓔ or ⊕ symbol on the Stream Configuration Overview and allows for editing existing or new streams.
At the top right corner, the page offers a drop-down list, where any of the existing streams can be selected. Notice that if you have changed contents of a stream without saving and then change to another stream, the changes will be lost without confirmation.
Besides all the existing streams, the selector has one last item, which allows for creating a new stream.

The page is divided into sections corresponding to fields in a frame, starting with Layer 2 fields (MAC addresses, possible VLAN tags, and EtherType) and followed by Layer 3 and Layer 4 fields.

Stream ID

Stream ID	Attached Clients	Stream Collection	Warnings
0	None	None	

Parameter	Description

Stream ID	The ID of the stream. The lower the ID, the higher precedence when matching. For example, if Stream #1 matches all frames with a multicast DMAC and Stream #2 matches a particular multicast DMAC, Stream #2 will never be hit. If editing an existing stream, this field is not editable, but it is if adding a new. If attempting to save a new stream with an ID already in use, a confirmation box will be shown.
Attached Clients	Streams don't do anything by themselves. They are utilized by PSFP and FRER for matching particular flows. This field indicates which of those two protocols - if any - are utilizing a particular stream. The number in parenthesis shows the PSFP or FRER instance the stream is used on. A stream may be part of a stream collection, in which case no clients can be attached directly to the stream itself. They must attach to the stream collection. In this case, these fields show a dash ("-"). See also next help topic. If no clients are attached to a stream, the stream will be installed into hardware anyway. This allows the user to create two streams, where the first e.g. matches a particular range of MAC addresses and the second matches the same range of MAC addresses plus some more. The second stream can be used in FRER or PSFP to match all the MAC addresses of the second stream, except that from the first.
Stream Collection	If a stream is part of a stream collection, this shows the ID of the stream collection. The ID is a hyperlink to the stream collection configuration page. If the stream is not part of a stream collection, the field just shows "None".
Warnings	Configuration of a stream may result in configurational warnings. For instance, if a stream is not instantiated on any ports, it is not of any use, and a warning will appear. A color indicates the warning state as follows: : The stream is not yet created : The stream has no configurational warnings : The stream has configurational warnings When yellow, hover the mouse over the image to see a list of configurational warnings. The possible warnings are as follows: • The stream does not have any member ports It is important to note that not all configurational errors are detected. An example of configurational errors that are not detected includes the case described under *Stream ID*, where stream #2 cannot be hit.

Besides the heading row, this table holds one row for Destination MAC (DMAC) address matching and another for Source MAC (SMAC) address matching.

MAC Addresses

DMAC/SMAC	Type	MAC	Mask
DMAC	Any	00:00:00:00:00:00	00:00:00:00:00:00
SMAC	Any	00:00:00:00:00:00	00:00:00:00:00:00

Parameter	Description
DMAC/SMAC	Indicates whether this row in the table is for DMAC or SMAC matching.
Type	The drop-down list for the DMAC offers the following options: **Any**: Match on any DMAC (default) **Multicast**: Match on multicast (but not broadcast) frames **Broadcast**: Match on broadcast frames **Unicast**: Match on unicast frames **Not Broadcast**: Match on unicast and multicast frames **Not Unicast**: Match on multicast and broadcast frames **MAC/Mask**: Match on a particular DMAC address with mask. The 'MAC' and 'Mask' fields will become available for editing when this is selected The drop-down list for the SMAC offers only the following options: **Any**: Match on any SMAC (default) **MAC/Mask**: Match on a particular SMAC address with mask. The 'MAC' and 'Mask' fields will become available for editing when this is selected
MAC	This field is only available when Type is set to 'MAC/Mask'. MAC address used in the matching along with the mask. See 'Mask' below for a description.

Mask	This field is only available when Type is set to 'MAC/Mask'. Mask to be used in the matching along with the MAC address. This cannot be all-zeros. Suppose you want to match on a particular OUI (the three most significant bytes of a MAC address), and suppose the OUI is **00:01:C1**. In that case, you will set 'MAC' to **00:01:C1:00:00:00** and 'Mask' to **FF:FF:FF:00:00:00**. When the configuration gets saved, bits in the 'MAC' address where the corresponding 'Mask' bits are zero will be cleared automatically.

Besides the heading row, this table holds one row for outer VLAN tag matching and another for inner VLAN tag matching.

VLAN Tags

Outer/Inner	Presence	Tag Type	VLAN	VLAN Mask	PCP	PCP Mask	DEI
Outer	Optional	Any	0	0x000	0	0x0	Any
Inner	Optional	Any	0	0x000	0	0x0	Any

Parameter	Description
Outer/Inner	Indicates whether this row in the table is for outer or inner VLAN tag matching.
Presence	This drop-down list controls whether the VLAN tag must be present (**Required**), must not be present (**Not Allowed**), or it doesn't matter whether the frame is VLAN tagged or not (**Optional**). If the Outer VLAN Tag's presence is set to Not Allowed, it is not possible to set the Inner VLAN Tag's presence to Required. Only when *Presence* is set to Required will the remaining fields of the row become available.
Tag Type	This field is only available when *Presence* is set to Required. If set to **C-tag**, only frames with EtherType (TPID) 0x8100 are matched. If set to **S-tag**, only frames with EtherTYpe (TPID) 0x88a8 are matched. If set to **Any**, both frames with EtherType (TPID) 0x8100 and 0x88a8 are matched.
VLAN	This field is only available when *Presence* is set to Required. VLAN ID used in the matching along with the VLAN mask. See *VLAN Mask* below for a description. Valid values are in the range **[0-4095]**.
VLAN Mask	This field is only available when *Presence* is set to Required. Mask to be used in the matching along with the VLAN ID. This is shown in hexadecimal notation but both hexadecimal and decimal values can be used when editing it. A mask of all-zeros means that any VLAN ID is matched. A mask of all-ones means that only the specified VLAN ID is matched. Any mask value in the range **[0x000-0xFFF]** is accepted. When the configuration gets saved, bits in the *VLAN* where the corresponding *VLAN Mask* bits are zero will be cleared automatically.
PCP	This field is only available when *Presence* is set to Required. PCP value used in the matching along with the PCP mask. See *PCP Mask* below for a description.
PCP Mask	This field is only available when *Presence* is set to Required. Mask to be used in the matching along with the PCP value. A mask of all-zeros means that any PCP value is matched. A mask of all-ones (0x7) means that only the specified PCP value is matched. When the configuration gets saved, bits in *PCP* where the corresponding *PCP Mask* bits are zero will be cleared automatically.
DEI	This field is only available when *Presence* is set to Required. If set to **0**, only frames with DEI = 0 are matched. If set to **1**, only frames with DEI = 1 are matched. If set to **Any**, the frame's DEI value is not used in the matching.

This section allows for configuring properties in frames beyond VLAN tags.
Use the Type drop-down list to select what properties to configure. The following are supported:

- **Any**: Don't use properties beyond VLAN tags for matching
- **EtherType**: Match on the frame's EtherType
- **LLC**: Match on LLC frames
- **SNAP**: Match on SNAP frames
- **IPv4**: Match on IPv4 frames
- **IPv6**: Match on IPv6 frames

Each of these types (except Any) have associated configurable parameters. These are outlined below per type.

Protocol

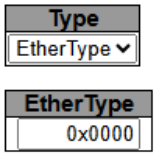

A hexadecimal (or decimal) value ranging from **0x600 to 0xffff** indicating the frame's EtherType, e.g. 0x800 for IPv4 (notice that IP has better matching properties than just the EtherType, so this was just an example).

Protocol

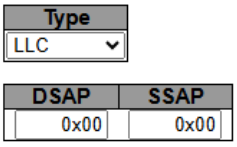

LLC (Logical Link Control) frames are frames with an EtherType/TypeLength field < **0x600**.
The LLC header carries two mandatory fields: DSAP and SSAP, whose values are used in the matching.
DSAP: The one-byte Destination Service Access Point to match.
SSAP: The one-byte Source Service Access Point to match.

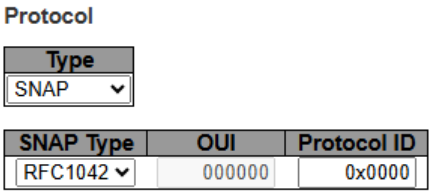

SNAP frames are LLC frames with **DSAP = 0xAA**, **SSAP = 0xAA** and control **field = 0x03**. The LLC fields are implicitly matched, while the SNAP protocol fields must be explicitly configured.

Parameter	Description
SNAP Type	This drop-down list holds the following selections: **RFC1042**: If selected, the OUI will be set to 00:00:00 **802.1H**: If selected, the OUI will be set to 00:00:F8 **Custom**: If selected, the OUI can be specified directly in the OUI field The primary purpose of providing this list of options is that people tend to forget often-used protocol OUIs. One may also simply select Custom and type e.g. the 802.1H OUI directly into the OUI field. But notice, that once this gets saved, the SNAP Type will change to 802.1H.
OUI	Only when the SNAP Type is set to Custom, will this field be editable. Type the OUI you with to match on as **six hexadecimal digits** but omit the colon. All six characters must be specified.
Protocol ID	The SNAP header's Protocol ID (a.k.a. PID) is a **16-bit integer** that can be entered as both a decimal and hexadecimal value. If *SNAP Type* is RFC1042, the Protocol ID must be an EtherType, that is, a value >= 0x600.

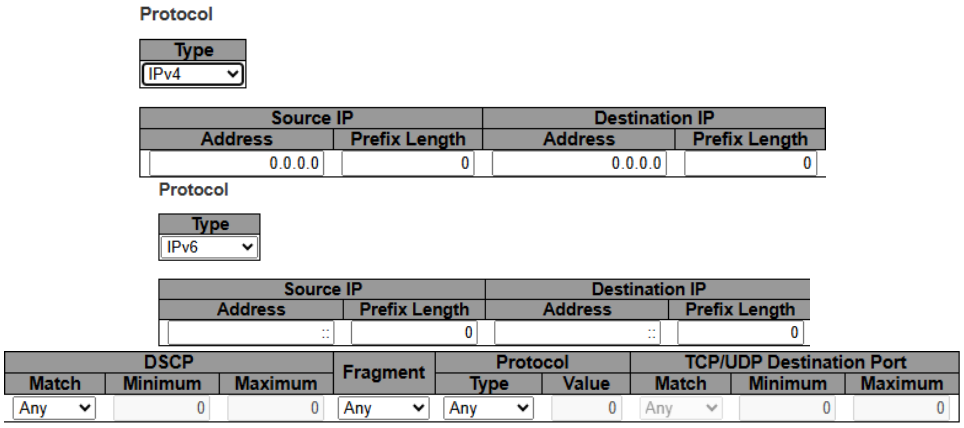

IPv4 frames are identified by hardware by having an EtherType of 0x800. IPv6 frames are identified by hardware by having an EtherType of 0x86dd. With the IPv4/IPv6 type, any IPv4/UPv6 frame can be matched or one or more fields from the IPv4/IPv6 header and (in case of UDP/TCP frames), also the UDP/TCP destination port can be matched.

Parameter	Description
Source IP Address	Specify an IPv4/IPv6 address in dotted decimal format to match against the frames' IPv4/IPv6 header's SIP field. Only if the Source IP Prefix Length is non-zero, will the frame be matched against a source IP address.
Source IP Prefix Length	Sets the number of bits in the Source IP address to match. **IPv4** If set to 0, no source IP matching will occur. If set to 32, all bits in the Source IP Address will be used. Otherwise, only the N first bits will be used in the matching, and the remaining bits will be cleared. The latter rule means that if you e.g. set the Source IP address to 1.2.3.4 and the Prefix Length to 16 and hit the Save button, the ".3.4" part will be cleared, so that the IP address will be shown as 1.2.0.0.

	IPv6 If set to 0, no source IP matching will occur. If set to 128, all bits in the Source IP Address will be used. Otherwise, only the N first bits will be used in the matching, and the remaining bits will be cleared. The latter rule means that if you e.g. set the Source IP address to 2001:db8::1 and the Prefix Length to 64 and hit the Save button, the ::1 part will be cleared, so that the IP address will be shown as 2001:db8:: Many tend to forget to set the prefix length after having specified an IP address. Therefore, if the prefix length is zero while the IP address is non-zero, a confirmation dialog will be shown asking whether this is really what the user wants.
Destination IP Address	This is where you will specify the IPv4/IPv6 header's DIP to match again. See *Source IP Address* above for details.
Destination IP Prefix Length	Sets the number of bits in the Destination IP address to match. See *Source IP Prefix Length* for details.
DSCP Match	One can choose to match on either **any DSCP value**, a **specific DSCP value**, or a **range of DSCP values**. Which selection is chosen determines the editability of the following two inputs.
DSCP Minimum	If *DSCP Match* is set to Value, this input determines the DSCP value to use in the matching. If *DSCP Match* is set to Range, this input determines the lower value in the range used in the matching. Valid values are **0-63**. If *DSCP Match* is set to Any, this input cannot be modified.
DSCP Maximum	If *DSCP Match* is set to Range, this input determines the higher value in the range used in the matching. Valid values are **0-63**. If *DSCP Match* is set to Any or Value, this input cannot be modified.
Fragment	If the IPv4 headers MF (More Fragments) bit is set, or the Fragment Offset is non-zero, the IPv4 frame is considered a fragment. This parameter can be set for IPv4 only. Set this drop-down list to **Any** to not using the fragment info for anything in the matching Set it to **No** to match IPv4 frames not considered fragments. Set it to **Yes** to match IPv4 frames considered fragments.
Protocol Type	The IPv4 header contains a one-byte Protocol field, which can be matched upon with this section. Set the drop-down list to Any to match any protocol field value. Set it to **TCP** to match on protocol field value 6. Set it to **UDP** to match on protocol field value 17. Set it to **Custom** to match on a custom protocol field value, which is input in the next field.
Protocol Value	If the *Protocol Type* is set to Custom, this one-byte field is used to input the protocol fiel value to match against. This may also be set to the value of TCP (6) or UDP (17). Doing so and saving, causes the Protocol Type to change to the relevant value.
TCP/UDP Destination Port Match	In case the *Protocol Type* is set to UDP or TCP, it is also possible to match on these two protocols' 16-bit destination port fields. Set the drop-down list value to **Any** to match all destination ports Set it to **Value** to match a particular destination port. Set it to **Range** to match a range of destination ports.
TCP/UDP Destination Port Minimum	If *TCP/UDP Destination Port Match* is set to Value, this input determines the destination port value to use in the matching. in the matching. If it is set to Range, this input determines the lower value in the range used in the matching. Valid values are **0-65535**. If the Match drop-down list is set to Any, this input cannot be modified.
TCP/UDP Destination Port Maximum	If *TCP/UDP Destination Port Match* is set to **Range**, this input determines the higher value in the range used in the matching. Valid values are **0-65535**. If the Match drop-down list is set to **Any** or **Value**, this input cannot be modified.

Port Members

Port	1	2	3	4	5	6	7	8	9	10	11	12
Member	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐

Check the ports on which this stream must be instantiated. By default, no ports are included.
Hover the mouse over a given port number of checkbox to see the port number as an interface name.

4.29.2 Stream Collections

A stream collection gathers one or more streams, which makes multiple streams available in functions such as FRER in generator mode, where different streams need to use the same sequence number generator.
TSN status information can be viewed at Monitor→TSN.

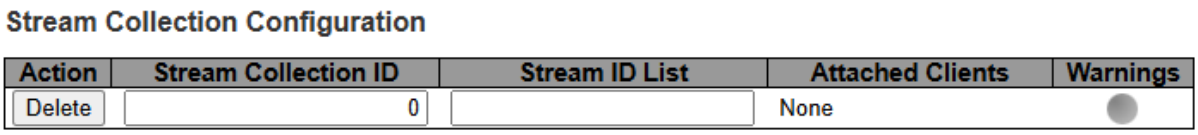

Parameter	Description
Delete	Click this button to delete a stream collection. The action happens immediately and without the user being prompted. If the stream collection was added with a click on the Add button, it is simply removed again.
Stream Collection ID	An ID uniquely identifying the stream collection. Valid values are in the range 1 to 63. This can only be changed if adding a new stream collection.
Stream ID List	List of stream IDs that make up this stream collection. individual stream IDs are separated by a comma, and ranges are separated by a hyphen. An example is **1,3,5,7-9**, where Stream IDs 1, 3, 5, 7, 8, and 9 are put into the stream collection. At most 8 streams can be stored in the same collection.
Attached Clients	Stream collections don't do anything by themselves. They are utilized by PSFP and FRER for matching particular flows. This field indicates which of those two protocols - if any - are utilizing a particular stream collection. The number in parenthesis shows the PSFP or FRER instance the stream collection is used on.
Warnings	Configuration of a stream collection may result in configurational warnings. For instance, if a stream collection doesn't carry any streams, it is not of any use, and a warning will appear. A color indicates the warning state as follows: : The stream collection is not yet created : The stream collection has no configurational warnings : The stream collection has configurational warnings When yellow, hover the mouse over the image to see a list of configurational warnings. The possible warnings are as follows: • No streams attached • No clients attached • At least one of the attached streams has configurational warnings

4.29.3 Frame Preemption

Frame Preemption is a TSN mechanism designed to reduce the latency of high-priority traffic by allowing a time-critical frame to interrupt the transmission of a non-critical frame. Without preemption, a high-priority frame must wait for a large "best-effort" frame to finish transmitting, creating a delay known as egress blocking or head-of-line blocking.

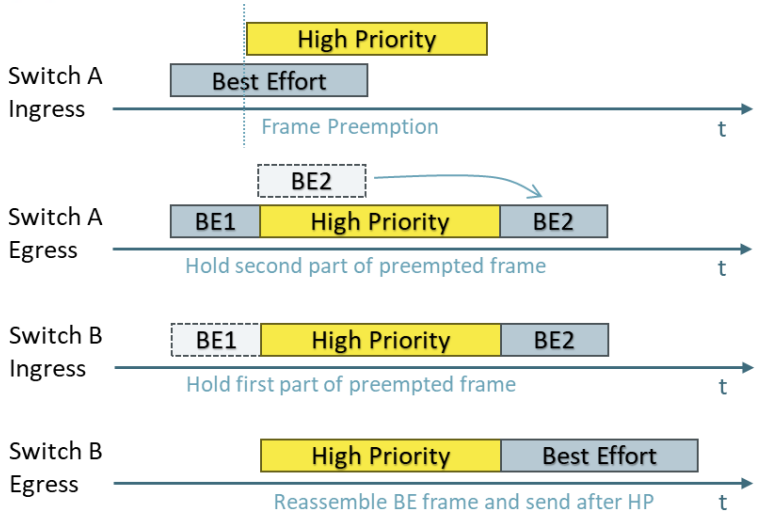

Switch A forwarding BE and HP frame to switch B

The Frame Preemption Status can be viewed at Monitor→TSN→Frame Preemption.

Frame Preemption Configuration

Port	Frame Preemption TX	Start without LLDP	Verify Disable TX	Preemptable Queues TX							
				Q0	Q1	Q2	Q3	Q4	Q5	Q6	Q7
*	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐
1	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐
2	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐
3	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐
4	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐
5	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐
6	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐
7	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐
8	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐
9	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐
10	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐
11	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐
12	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐

Parameter	Description
Port	The logical port for the settings contained in the same row.
Frame Preemption TX	The value of the 802.3br aMACMergeEnableTx parameter for the port. This value determines whether frame preemption is enabled (TRUE) or disabled (FALSE) in the MAC Merge sublayer in the transmit direction.
Start without LLDP	When this field is checked, Frame Preemption will be active when Frame Preemption TX is checked.
Verify Disable TX	The value of the 802.3br aMACMergeVerifyDisableTx parameter for the port. This value determines whether the verify function is disabled (TRUE) or enabled (FALSE) in the MAC Merge sublayer in the transmit direction.
Preemptable Queues	This parameter is the administrative value of the preemption status for the priority. If checked, it takes value preemptable if frames queued for the priority are to be transmitted using the preemptable service for the Port. If not checked, it takes value express if frames queued for the priority are to be transmitted using the express service for the Port and preemption is enabled for the Port.

4.29.4 TAS

The Time-Aware Scheduler (TAS) defined in IEEE 802.1Qbv uses a Gate Control List (GCL) to manage egress traffic by opening and closing gates for specific priority queues based on a precise, synchronized schedule. This mechanism creates dedicated time slots for critical traffic, ensuring it can traverse the switch without interference from lower-priority data. By strictly isolating traffic classes in the time domain, TAS provides the deterministic latency and zero-jitter delivery required for industrial control loops.
The TAS status can be viewed at Monitor→TSN→TAS.

4.29.4.1 Ports

This page shows the current TAS configuration per port and allows the user to change them.

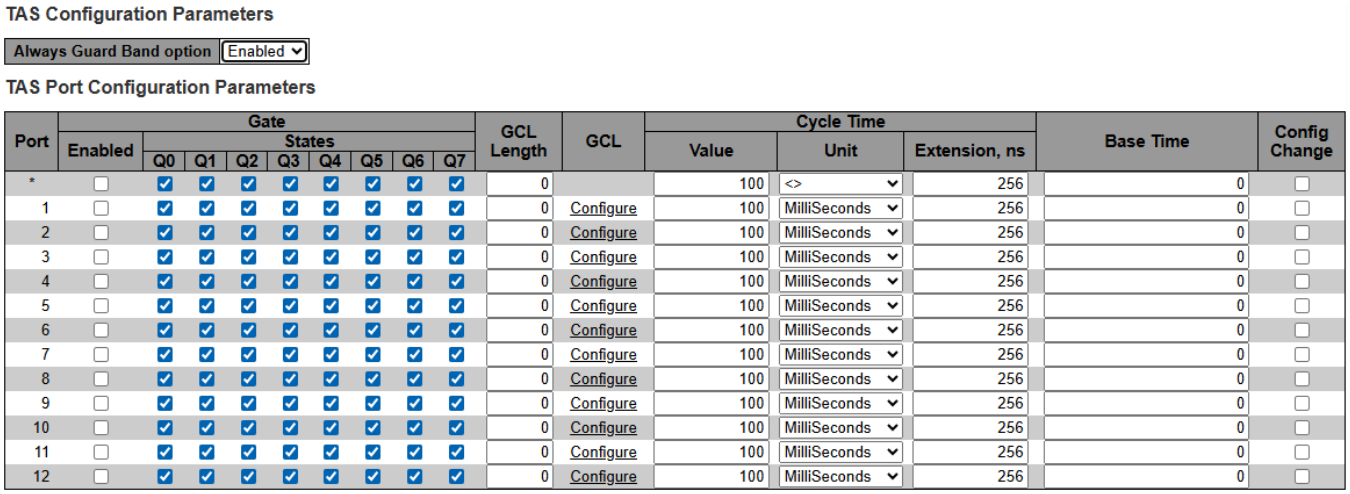

Parameter	Description
Always Guard Band	The Always Guard Band option defines how the guard band values are calculated. If a Gate Control List do not contain SetAndHold and/or SetAndRelease operations the Always Guard Band option has no effect. If a Gate Control List do contain SetAndHold and SetAndRelease operations then: • When Always Guard Band is Enabled, a guard band is implemented on all queues, both Express and Preemptible queues. • When Always Guard Band is Disabled, a guard band is only implemented on Preemptible queues.
Port	The logical port for the settings contained in the same row.
Gate Enabled	The Enabled parameter determines whether traffic scheduling is active (true) or inactive (false).
Gate States	The initial value of the port open states that is used when no Gate Control List is active on the Port.
GCL Length	The Admin Gate Control List length parameter for the Port. Valid range is **`0-256`**. The integer value indicates the number of entries Gate Control Elements in the Gate Control List. If you change the value, press the Save button before configuring the Gate Control List by pressing the GCL link.
GCL	A link to the Gate Control List parameter configuration.
Cycle Time	The Admin value of the gating cycle for the Port. The Admin Cycle Time variable is a rational number of seconds, defined by value and a unit.
Cycle Time Value	The Admin Cycle Time is defined by this number of units defined in the Unit field. The Admin Cycle Time is a value in the range 1-999999999, and combined with the Cycle Time Unit the value shall be in the range **`256-999999999 nanoseconds`**. The default value is 100 milliseconds.
Cycle Time Unit	The Admin Cycle Time unit. May be **`milliseconds`**, **`microseconds`** or **`nanoseconds`**.

Cycle Time Extension	An integer number of nanoseconds in the range **256-999999999**, defining the maximum amount of time by which the gating cycle for the Port is permitted to be extended when a new cycle configuration is installed. The default value is 256 nanoseconds.
Base Time	The Admin value of base time, expressed as an IEEE 1588 precision time protocol (PTP) timescale.
Config Change	The Configuration Change parameter signals the start of a configuration change. After a successful configuration change, the configured Admin values will become the Oper values, which are displayed in the Monitor/TSN/TAS web page If the value of parameter Base Time is in the future, the configuration change will be executed at Base Time. If Base Time is in the past, the configuration change will be executed as soon as possible. In practice it will be within approximately 2 seconds, at a time which is an integral number of Cycle Time ahead of the configured value of Base Time. This way, the synchronization between schedules in elements across a scheduled network can be maintained.

4.29.4.2 Max SDU

The Maximum Service Data Unit (MSDU) defines the largest payload size permitted for a specific traffic class or stream within the TSN network. The Maximum SDU size parameter is used to calculate the guard band time.

$$T_{gb} = Maximum\ SDU * 8 / LINK\_SPEED\ (sec)$$

If frame preemption is enabled and a gate operation is SetAndHold, the guard band time in preemptable queues is automatically selected as the frame preemption minimum fragment size plus 64 bytes.
A queue is said to be preemptible, if frame preemption is enabled, and if this queue is not opened in a SetAndHold gate operation.

TAS SDU Configuration

Port	Max SDU Size							
	Q0	Q1	Q2	Q3	Q4	Q5	Q6	Q7
*	1536	1536	1536	1536	1536	1536	1536	1536
1	1536	1536	1536	1536	1536	1536	1536	1536
2	1536	1536	1536	1536	1536	1536	1536	1536
3	1536	1536	1536	1536	1536	1536	1536	1536
4	1536	1536	1536	1536	1536	1536	1536	1536
5	1536	1536	1536	1536	1536	1536	1536	1536
6	1536	1536	1536	1536	1536	1536	1536	1536
7	1536	1536	1536	1536	1536	1536	1536	1536
8	1536	1536	1536	1536	1536	1536	1536	1536
9	1536	1536	1536	1536	1536	1536	1536	1536
10	1536	1536	1536	1536	1536	1536	1536	1536
11	1536	1536	1536	1536	1536	1536	1536	1536
12	1536	1536	1536	1536	1536	1536	1536	1536

Parameter	Description
Port	The logical port for the settings contained in the same row.
Maximum SDU	The value of the Maximum SDU size parameter for the traffic class supported by the port. This value is represented as an unsigned integer in the range 0-10240. A value of 0 is interpreted as the Maximum SDU size supported by the underlying MAC: 10240. The default value of the Maximum SDU parameter is 1536.

4.29.5 PSFP

Per Stream Filtering and Policing (PSFP) functions allow filtering and policing decisions, and subsequent frame queuing decisions on a per-stream basis. PSFP is supported by a table of stream filters that determine the filtering and policing actions that are to be applied to frames received on ingress ports.
PSFP status and information can be viewed at Monitor→TSN→PSFP.

4.29.5.1 Flow Meters

The PSFP Flow Meter monitors the data rate of individual streams to ensure they do not exceed their pre-allocated bandwidth limits. It utilizes a "token bucket" algorithm to measure both average bit rates and burst sizes, protecting the network from "babbling idiot" nodes or misconfigured devices. If a stream violates its traffic contract, the meter can be configured to drop the non-compliant frames or remark their priority to prevent congestion in downstream queues.

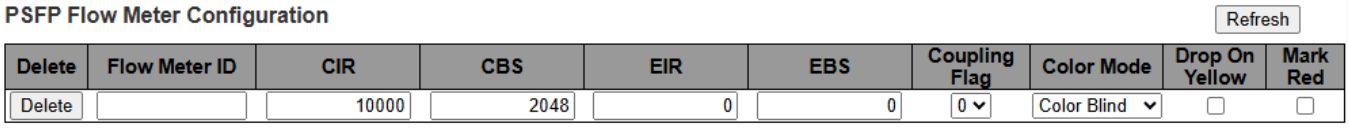

Parameter	Description
Delete	Check to delete the flow meter instance. It will be deleted right away without user confirmation.
Flow Meter ID	ID of flow meter instance. When adding a new, this can be set.
CIR	Committed Information Rate measured in kbps. The value will be adjusted to the closest value supported by hardware.
CBS	Committed Burst Size measured in bytes. The value will be adjusted to the closest value supported by hardware.
EIR	Excess Information Rate measured in kbps. The value will be adjusted to the closest value supported by hardware.
EBS	Excess Burst Size measured in bytes. The value will be adjusted to the closest value supported by hardware.
Coupling Flag	When '1', frames that would overflow the committed bucket will be added to the excess bucket unless it's full. Mostly relevant when the Color Mode is Color Aware.
Color Mode	If set to Color Blind, each frame starts green and is then marked according to the policer's operation. If set to Color Aware, the frame starts at the classified color (green or yellow), and is then marked according to the policer's operation.
Drop On Yellow	If checked, frames marked yellow are discarded. If unchecked, frames will have their DEI value set to 1.
Mark Red	If checked, all subsequent frames are discarded if a red frame is seen.

4.29.5.2 Stream Gates

The Stream Gate is a key component of Per-Stream Filtering and Policing that controls the flow of individual streams based on a time-synchronized schedule. Similar to the Time-Aware Scheduler, it uses a Gate Control List to transition between Open and Closed states, ensuring that frames only enter the switch fabric during their authorized time windows. This prevents rogue or malfunctioning devices from injecting traffic into the network outside of their allocated schedule, thereby protecting the bandwidth of other critical streams.

The PSFP Stream Gate status can be viewed at Monitor→TSN→PSFP→Stream Gate Status.

This page allows the user to inspect the current PSFP stream gate configurations, along with options to add, edit or delete such configurations.

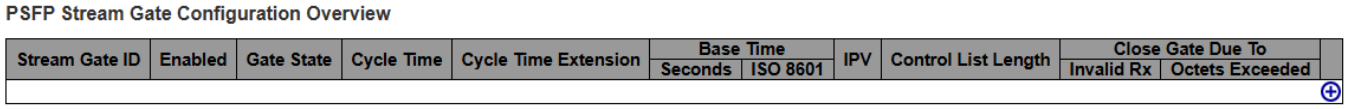

This page allows the user to inspect and modify current PSFP stream gate configurations, along with options to create new.
At the top right corner, the page offers a drop-down list, where any of the existing stream gates can be selected. Notice that if you have changed contents of a stream gate without saving and then change to another stream gate, the changes will be lost without confirmation.
The page is divided into three sections, as outlined below.

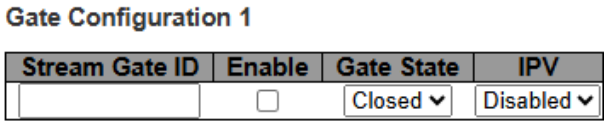

This part contains parameters that take effect immediately after the configuration is saved - whether or not the next section's "Config Change" checkbox is checked.

Parameter	Description
Stream Gate ID	The ID of this gate instance. If editing an existing stream gate, this field is not editable, but it is if adding a new. If attempting to save a new stream gate with an ID already in use, a confirmation box will be shown.
Enable	If checked, the stream gate is active in hardware. Otherwise, it is inactive.
Gate State	The gate's initial state.
IPV	Select the frames initial priority value (egress queue). It may be overridden by a control list entry later. If set to Disabled, the frame's classified priority will be used to point out the egress queue - also unless overridden by a control list entry later).

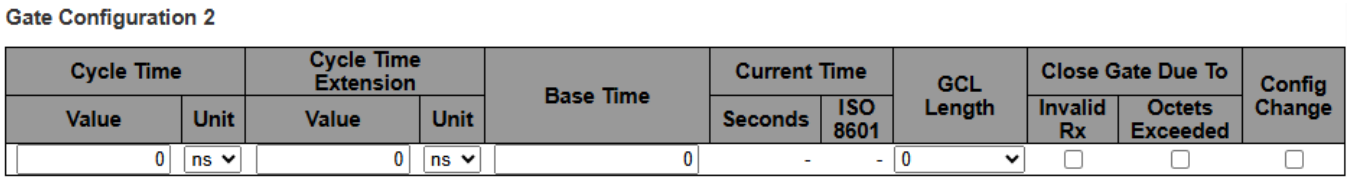

This section configures properties of the stream gate that only take effect when the gate is enabled and the "Config Change" checkbox is checked. Whether the configuration becomes pending or applied directly to hardware depends on whether the configured Base Time is in the past (takes effect immediately) or in the future (becomes pending and only takes effect when the Base Time is reached).

Parameter	Description
Cycle Time	Set the gate's cycle time. A cycle time of up to 1 second can be specified. The Cycle Time consists of a value and a unit (milliseconds, microseconds, or nanoseconds). The cycle time cannot be set lower than the sum of all the enabled gate control list entries' Time Interval.
Cycle Time Extension	Set the gate's cycle time extension. Up to 1 second can be specified, in the same manner as the Cycle Time.
Base Time	The administrative value of the Base Time parameter for the gate. The value must be specified in seconds an optional dot (.) and up to 9 digits indicating the sub seconds. The current time can be seen in the next field (Current Time) if you are not creating a new stream gate. This can be used to judge what to set the Base Time to. If setting it to a time in the past, the configuration will take effect immediately after you have checked the "Config Change" checkbox and clicked Save. Otherwise, the configuration will become pending and only take effect when Current Time reached Base Time.
Current Time	To assist in judging what to set the Base Time to or judge when an existing stream gate configuration will take effect, the current time is displayed in two formats: The "Seconds" value indicates the time in seconds and sub seconds since epoch. The "ISO 8601" value represents the same, but is displayed as a string formatted according to ISO 8601.

GCL Length	Length of the gate control list. The number of enabled items in the "Gate Control List Configuration" section depends on the chosen gate control list length.
Close Gate Due To Invalid Rx	If checked, the stream gate gets permanently closed if receiving a frame during a closed gate state.
Close Gate Due To Octets Exceeded	If checked, the stream gate gets permanently closed if receiving a frame that exceeds a gate control list entry's "Octet Max".
Config Change	This is the one-shot parameter that - when checked and the Save button is clicked - causes the current configuration to take effect or become pending. This will always be unchecked when the page gets refreshed.

Gate Control List Configuration

Index	Gate State	IPV	Time Interval		Octet Max
			Value	Unit	
0	Closed	Disabled	1	ns	0
1	Closed	Disabled	1	ns	0
2	Closed	Disabled	1	ns	0
3	Closed	Disabled	1	ns	0

This section allows for configuring the individual gate control list entries. The number of editable items depends on the selected value of "GCL Length" in the previous section.
Changes will only take effect (or become pending) when the previous section's Config Change is checked and the Save button is clicked.

Parameter	Description
Index	The index of the gate control entry. These are zero-based.
Gate State	Selects the gate's state during this entry's active period.
IPV	Use this (Internal Priority Value) to force frames arriving during this gate's active period to a particular egress queue. Set to "Disabled" to let the frame keep its current IPV.
Time Interval	Select this control list entry's active period. The value cannot exceed 999,999,999 nanoseconds and the sum of the time interval for all enabled gate control list entries cannot exceed the gate's cycle time.
Octet Max	Select the maximum frame size that may arrive in this entry's active period. Set to 0 to allow any frame size. If a frame's size exceeds the configured maximum size, it gets discarded. On top of this, if the "Close Gate Due To Octets Exceeded" is checked, subsequent frames also get discarded until administratively opened again (on the PSFP Stream Gate's Status page).

4.29.5.3 Stream Filters

PSFP Stream Filter status and statistics can be viewed at Monitor→TSN→PSFP→Stream Filter Status and Monitor→TSN→PSFP→Stream Filter Statistics.

PSFP Stream Filter Configuration

Refresh

Delete	Stream Filter ID	Stream Type	Stream ID	Stream Collection ID	Flow Meter Enable	Flow Meter ID	Stream Gate Enable	Stream Gate ID	Maximum SDU Size	Block Oversize Frame Enable	Warnings
Delete		Stream							0		

Parameter	Description
Delete	Click to delete the stream filter instance. It will be deleted right away without any request for confirmation. If the stream filter was added with a click on the Add button, it is simply removed again.
Stream Filter ID	An ID uniquely identifying the stream filter. Valid values are in the range 0 to 254. This can only be changed if adding a new stream filter.

Stream Type	PSFP allows for two different methods for specifying streams that should map to this PSFP stream filter instance. The first, "Stream", allows for setting one particular stream ID. The second, "Stream Collection", allows for specifying an ID of a stream collection. A stream collection compiles multiple streams into one single. The two methods are mutually exclusive.
Stream ID	Specify the ID of the stream that should map to this PSFP stream filter. This field is only available if "Stream Type" is set to "Stream ID".
Stream Collection ID	Specify the ID of a stream collection that should map to this PSFP stream filter. This field is only available if "Stream Type" is set to "Stream Collection".
Flow Meter Enable	Check this to enable setting a Flow Meter ID that this stream filter will use.
Flow Meter ID	The ID of the Flow Meter that this stream filter will use. This field is only available if "Flow meter Enable" is checked.
Stream Gate Enable	Check this to enable setting a Stream Gate ID that this stream filter will use.
Stream Gate ID	The ID of the Stream Gate that this stream filter will use. This field is only available if "Stream Gate Enable" is checked.
Maximum SDU Size	The maximum allowed frame size for the filter. Any frame exceeding this value will be discarded. A value of 0 disables this feature.
Block Oversize Frame Enable	Whenever a frame gets discarded because its SDU size is greater than the configured Maximum SDU Size, this one control what shall happen with subsequent frames that go through the stream filter. If checked, subsequent frames will also be discarded whether they are larger or smaller than the configured Maximum SDU Size. Otherwise, they will remain being subject to only the Maximum SDU Size check - if enabled. An administrative action is required to reset the discarding. Use the PSFP Stream Filter Status page to perform this action.
Warnings	Configuration of a stream filter may result in configurational warnings. For instance, if a stream filter is neither assigned to a Flow Meter nor a Stream Gate, it is not of any use, and a warning will appear. A color indicates the warning state as follows:  : The stream filter is not yet created  : The stream filter has no configurational warnings : The stream filter has at least one configurational warning When yellow, hover the mouse over the image to see a list of configurational warnings. The possible warnings are as follows: • Neither a stream or a stream collection is specified • The specified stream ID does not exist • The specified stream collection ID does not exist • Unable to attach to the specified stream, possibly because it is part of a stream collection • Unable to attach to the specified stream collection • The specified stream has configurational warnings. • The specified stream collection has configurational warnings. • The specified flow meter ID does not exist • The specified stream gate ID does not exist • The specified stream gate is not enabled

4.29.6 FRER

Frame Replication and Elimination for Reliability (FRER) is specified by the IEEE 802.1CB-2017 standard and provides increased reliability (reduced packet loss rates) for a stream.
This is achieved by:

1. Sequence numbering and replicating packets in the talker (source) end system and/or in relay systems in the network.
2. Eliminating those replicates in the listener (destination) end system and/or in other relay systems.

FRER status and statistics can be viewed at Monitor→TSN→FRER→Status and Monitor→TSN→FRER→Statistics.

At the top right corner, the page offers a drop-down list, where any of the existing FRER instances can be selected. Notice that if you have changed contents of a FRER instance without saving and then change to another instance, the changes will be lost without confirmation.
Besides all the existing FRER instances, the selector has one last item, which allows for creating a new instance.

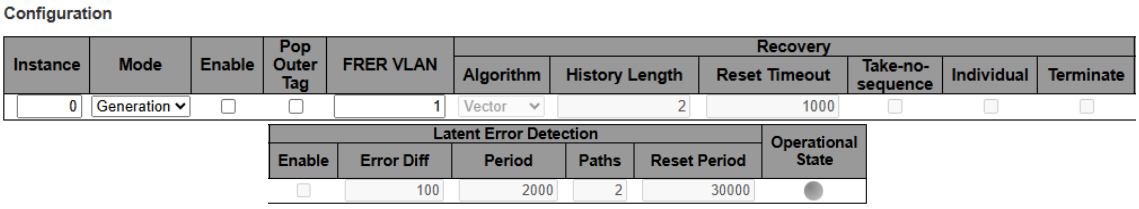

Parameter	Description
Mode	Decides if this FRER instance shall run in generation or recovery mode. Default is generation.
Enable	Enable or disable a FRER instance.
Pop Outer Tag	Only available in generation mode. When checked, a possible outer VLAN tag in the ingress frames gets popped before egressing with an R-tag. When unchecked, a possible outer VLAN tag in the ingress frames is preserved beneath the R-tag on egress.
FRER VLAN	Select the VLAN ID that ingress flows get classified to.
Recovery Algorithm	IEEE 802.1CB-2017 requires implementations to provide two different recovery function algorithms, match and vector. **Match** is the simplest algorithm: It basically says: Discard all packets with a sequence number equal to the last sequence number seen. Accept all others. The algorithm also comes with a reset timer that - when it expires - causes the algorithm to accept any sequence number - even the same as the previous. The reset timer is restarted every time a packet is accepted. The match algorithm counts the number of times the reset timer has expired and the number of passed, discarded, and out-of-order packets. Out-of-order happens when the sequence number of a given packet is not one higher than the previous (and the timer has not expired). **Vector** is somewhat more complicated. When a packet with a given sequence number arrives, it must be within the previous accepted packet's sequence number +/- a configurable history length, or it will be discarded. If the packet is already seen (within the history length window), it is also discarded. Also this algorithm comes with a reset timer that - when it expires - causes the algorithm to accept any sequence number next time a packet arrives. The reset timer is restarted every time a packet is accepted. The Vector algorithm counts the number of times the reset timer has expired and the number of passed, discarded, out-of-order, and so-called rogue packets. Out-of-order happens when the sequence number of a given packet is "older" than a previous packet's (taking wrap-around into account), and the packet has not been accepted before. Out-of-order packets are accepted. Rogue packets are packets with a sequence number beyond the history length window. Rogue packets are also counted as discarded. Furthermore, the Vector algorithm counts lost packets, that is, the number of unreceived sequence numbers when the history window moves. Both algorithms also count the number of packets arriving without an R-tag. This is done with the tagless counter. By default, such packets will be discarded. A per-FRER instance parameter recovery take-no-sequence, however, allows such frames to pass through. *Note: The 802.1CB standard utilizes the frerSeqRcvyTakeNoSequence only in the vector algorithm, but the switches that the present guide is meant for also utilize it in the match algorithm.* *Note: This feature should only be used on terminating switches, because such tagless packets will be R-tagged (with sequence number 0) on their way out on non-terminating switches.*

	The selected algorithm on a given FRER instance will be used in both compound and individual recovery functions. Default is the Vector algorithm.
Recovery History Length	Configure the recovery functions history length. Valid range is **2-32** and default is 2.
Recovery Reset Timeout	Configure recovery function's reset timeout in milliseconds. Valid range is **1-4095** and default is 1000.
Recovery Take-no-sequence	Select this option to accept all frames whether they are R-tagged or not.
Recovery Individual	Individual recovery means that a member stream undergoes recovery before it reaches the compound recovery function. The compound recovery function sits on each and every egress port in the FRER instance. The one and only thing that individual recovery can do that compound recovery can not is to filter out member streams that keep presenting the same R-tag sequence number because of a defect transmitter. It goes like this: Suppose the transmitter of member stream 1 is working perfectly. It will send out frames with an increasing sequence number and wrap back to 0 after 65535 frames. Suppose the transmitter of member stream 2 is sending out the same frame with the same sequence number, X, over and over again. If we only had a compound recovery function, that function would at times be presented with frames with sequence number X from stream 1 and sequence number X from stream 2, and the first of these two frames would be sent to the egress port. So - depending on timing - sometimes the frame with sequence number X would come from stream 1 and sometimes it would come from the erroneous stream 2. The effect of enabling individual recovery is to have the individual recovery function for stream 2 filter out all identically numbered frames before they are presented to the compound recovery function. This is a very unlikely situation, and most network administrators will not need individual recovery. Moreover, individual recovery is very expensive in terms of hardware resources: Every ingress stream needs an individual recovery function per egress port. So if a FRER instance defines 8 ingress streams and 8 egress ports, the switch needs 64 individual recovery instances - just for this one FRER instance.
Recovery Terminate	Select this option to strip an R-Tag from a frame before presenting it on egress.
Latent Error Detection	The purpose of latent error detection is to raise a flag if the number of discarded packets is "relatively few" compared to the number of passed packets. The algorithm relies on four user inputs: Period, Reset period, Paths, Error difference. The reset function algorithm is as follows: Every Reset period milliseconds, read number of passed and discarded packet counters and set a per-FRER instance variable, CurDiff, as follows: **CurDiff = passed_packets * (paths - 1) - discarded_packets** The test function algorithm is as follows: Every timeout milliseconds, read the discarded and passed packet counters, and perform the following: **diff = Abs(CurDiff - (passed_packets * (paths - 1) - discarded_packets))** **if (diff > difference) {** ** raise_flag();** **}** Basically, it says: If you expect N member streams to ingress this FRER instance, N-1 of these member streams are expected to be discarded, and only one is expected to pass. To allow for some slack due to random packet losses and the fact that counters are not necessarily read simultaneously, set the difference to account for that. The reset function makes sure that CurDiff is updated to avoid that occasional packet losses do not accumulate forever.
Latent Error Detection Enable	Enable/disable Latent error detection.

Latent Error Detection Error Diff	The number of packets "allowed" to be in difference without raising the flag. Valid range is **`0-10000000`** and default is 100.
Latent Error Detection Period	The number of milliseconds between invoking the test function. Valid range is **`1000-86400000`** and default is 2000.
Latent Error Detection Paths	The number of member streams expected to ingress this FRER instance. Valid range is **`2-8`** and default is 2.
Latent Error Detection Reset Period	The number of milliseconds between invoking the reset function. Valid range is **`1000-86400000`** and default is 30000.
Operational State	This shows the current operational state of a FRER instance. A color indicates the state as follows: : The instance is not created or not enabled : The instance is enabled with no configurational warnings : The instance is enabled, but there are configurational warnings : The instance is enabled, but an internal error has occurred. See console or crashlog for details When yellow, hover the mouse over the image to see a list of configurational warnings. The possible warnings are as follows: • At least one of the ingress streams doesn't exist • Failed to attach to at least one of the ingress streams, possibly because it is part of a stream collection • At least one of the ingress streams has configurational warnings • There is an overlap between ingress and egress ports • In generation mode, only one egress port is specified • At least one of the ingress ports doesn't have link • At least one of the egress ports doesn't have link • At least one of the egress ports is not member of the FRER VLAN • At least one of the egress ports is blocked by STP • At least one of the egress ports is blocked by MSTP • Stream collection doesn't exist • Failed to attach to stream collection • The specified stream collection has configurational warnings It is highly recommended to have a green color on enabled instances. It is important to notice that not all configuration errors can be detected by the FRER software.

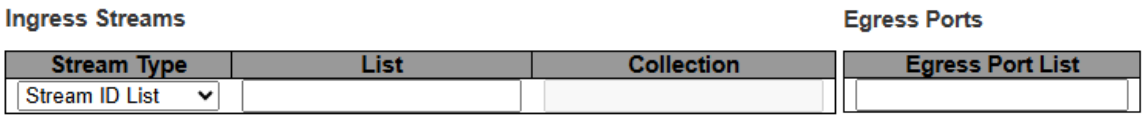

Parameter	Description
Stream Type	Two different methods for specifying streams that should map to this FRER instance exist. The first, **`Stream ID List`**, allows for setting up to 8 ingress streams in recovery mode. In generator mode, only one ingress stream can be specified. If you need more, use the second option. The second, **`Stream Collection`**, allows for specifying an ID of a stream collection. A stream collection compiles multiple streams into one single, which can be used in both generator and recovery mode, except if individual recovery is enabled. The two methods are mutually exclusive.
Stream ID List	Specify the ingress streams that should map to this FRER instance. Only one stream ID can be specified in generator mode. If more are needed, use an ingress stream collection. In other modes, up to 8 ingress streams can be specified. This field is only available if *Stream Type* is set to `Stream ID List`.
Stream Collection	Specify the ID of an ingress stream collection that should map to this FRER instance. This field is only available if *Stream Type* is set to `Stream Collection`.
Egress port list	Specify the egress ports that this FRER instance will hit. There is a maximum of 8 egress ports per FRER instance.

4.30 Mirroring

Port Mirroring is a troubleshooting feature and can be used by an administrator to analyze data traffic and debug network problems. The selected traffic can be mirrored or copied on a destination port where a network analyzer can be attached to analyze the network traffic.

Remote Mirroring is an extend function of Mirroring. It can extend the destination port to another switch. So the administrator can analyze the network traffic on the other switches.

If you want to get the tagged mirrored traffic, you have to set VLAN egress tagging as "Tag All" on the reflector port. On the other hand, if you want to get untagged mirrored traffic, you have to set VLAN egress tagging as "Untag ALL" on the reflector port.

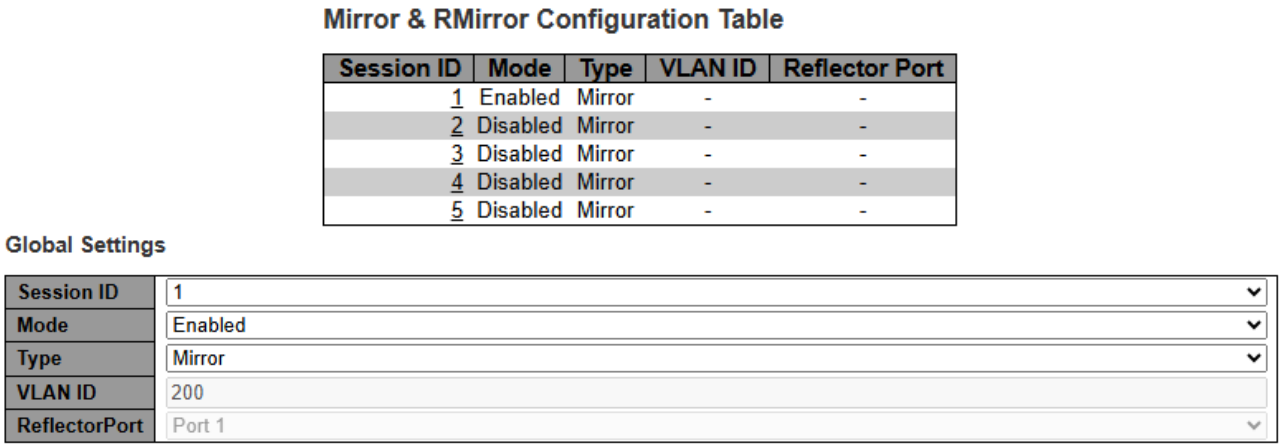

Mirror & RMirror Configuration Table

Session ID	Mode	Type	VLAN ID	Reflector Port
1	Enabled	Mirror	-	-
2	Disabled	Mirror	-	-
3	Disabled	Mirror	-	-
4	Disabled	Mirror	-	-
5	Disabled	Mirror	-	-

Global Settings

Session ID	1
Mode	Enabled
Type	Mirror
VLAN ID	200
ReflectorPort	Port 1

Parameter	Description
Session ID	Select session id to configure.
Mode	To Enabled/Disabled the mirror or Remote Mirroring function.
Type	Select switch type. **`Mirror`**: The switch is running on mirror mode. The source port(s) and destination port are located on this switch. **`RMirror source`**: The switch is a source node for monitor flow. The source port(s), reflector port are located on this switch. **`RMirror destination`**: The switch is an end node for monitor flow. The destination port(s) is located on this switch.
VLAN ID	The VLAN ID points out where the monitor packet will copy to. The default VLAN ID is 200.
Reflector Port	The reflector port is a method to redirect the traffic to Remote Mirroring VLAN. Any device connected to a port set as a reflector port loses connectivity until the Remote Mirroring is disabled. In the stacking mode, you need to select switch ID to select the correct device. If you shut down a port, it cannot be a candidate for reflector port. If you shut down the port which is a reflector port, the remote mirror function cannot work. *Note:* • *The reflector port needs to select only on Source switch type.* • *The reflector port needs to disable MAC Table learning and STP.* • *The reflector port only supports on pure copper ports.*

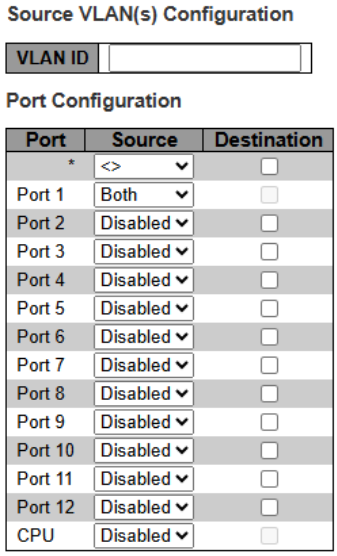

Parameter	Description
Source VLAN ID	The switch supports VLAN-based Mirroring. If you want to monitor some VLANs on the switch, you can set the selected VLANs in this field. *Note: The Mirroring session shall have either ports or VLANs as sources, but not both.*
Port	The logical port for the settings contained in the same row.
Source	Select mirror mode. **Disabled** Neither frames transmitted nor frames received are mirrored. **Both** Frames received and frames transmitted are mirrored on the Destination port. **Rx only** Frames received on this port are mirrored on the Destination port. Frames transmitted are not mirrored. **Tx only** Frames transmitted on this port are mirrored on the Destination port. Frames received are not mirrored.
Destination	Select destination port. This checkbox is designed for mirror or Remote Mirroring. The destination port is a switched port that you receive a copy of traffic from the source port. Note: On mirror mode, the device only supports one destination port. The destination port needs to disable MAC Table learning.

Configuration Guideline for All Features

When the switch is running on Remote Mirroring mode, the administrator also needs to check whether or not other features are enabled or disabled. For example, if the administrator does not disable MSTP on the reflector port, all monitor traffic will be blocked on the reflector port.
All recommended settings are described as follows.

	Impact	Source Port	Reflector Port	Intermediate Port	Destination Port	RMirroring VLAN
ARP Inspection	High	---	* disabled	* disabled	---	---
ACL	Critical	---	* disabled	* disabled	* disabled	---
DHCP Relay	High	---	* disabled	* disabled	---	---
DHCP Snooping	High	---	* disabled	* disabled	---	---
IP Source Guard	Critical	---	* disabled	* disabled	* disabled	---
IPMC/IGMP-SNP	Critical	---	---	---	---	un-conflict
IPMC/MLD-SNP	Critical	---	---	---	---	un-conflict
LACP	Low	---	---	---	o disabled	---

LLDP	Low	---	---	---	o disabled	---
MAC Learning	Critical	---	* disabled	* disabled	* disabled	---
MSTP	Critical	---	* disabled	---	o disabled	---
MVR	Critical	---	---	---	---	un-conflict
NAS	Critical	---	* authorized	* authorized	* authorized	---
PSEC	Critical	---	* disabled	* disabled	* disabled	---
QoS	Critical	---	* unlimited	* unlimited	* unlimited	---
UPnP	Low	---	---	---	o disabled	---
Mac-based VLAN	Critical	---	* disabled	* disabled	---	---
Protocol-based VLAN	Critical	---	* disabled	* disabled	---	---
vlan_translation	Critical	---	* disabled	* disabled	* disabled	---
voice_vlan	Critical	---	* disabled	* disabled	---	---
mrp	Low	---	---	---	o disabled	---
mvrp	Low	---	---	---	o disabled	---

Note:
***** : Strongly recommended
o : Optional
Critical impact: Prevents a significant number of packets from being mirrored
High impact: Prevents some packets from being mirrored
Low impact: Adds some additional packets in the mirrored traffic

4.31 UPnP

The goals of Universal Plug and Play (UPnP) are to allow devices to connect seamlessly and to simplify the implementation of networks in the home (data sharing, communications, and entertainment) and in corporate environments for simplified installation of computer components.

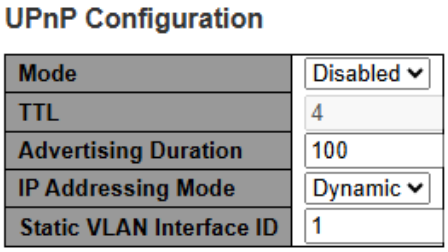

Parameter	Description
Mode	Indicates the UPnP operation mode. Possible modes are: **`Enabled`**: Enable UPnP mode operation. **`Disabled`**: Disable UPnP mode operation. When the mode is enabled, two ACEs are added automatically to trap UPNP related packets to CPU. The ACEs are automatically removed when the mode is disabled.
TTL	The TTL value is used by UPnP to send SSDP advertisement messages. This value can not be changed.
Advertising Duration	The duration, carried in SSDP packets, is used to inform a control point or control points how often it or they should receive an SSDP advertisement message from this switch. If a control point does not receive any message within the duration, it will think that the switch no longer exists. Due to the unreliable nature of UDP, in the standard it is recommended that such refreshing of advertisements to be done at less than one-half of the advertising duration. In the implementation, the switch sends SSDP messages periodically at the interval one-half of the advertising duration minus 30 seconds. Valid values are in the range **`100 to 86400`**. Specified in seconds.

IP Addressing Mode	IP addressing mode provides two ways to determine IP address assignment: **Dynamic**: Default selection for UPnP. UPnP module helps users choose the IP address of the switch device. It finds the first available system IP address. **Static**: User specifies the IP interface VLAN for choosing the IP address of the switch device.
Static VLAN Interface ID	The index of the specific IP VLAN interface. It will only be applied when IP Addressing Mode is static. Valid configurable values range from **1 to 4095**. Default value is 1.

4.32 PTP

The Precision Time Protocol (PTP) provides the nanosecond-level synchronization required for TSN features to function across a distributed network. It operates by electing a Grandmaster clock that distributes a common time reference to all nodes via specialized synchronization packets and path delay measurements.
PTP status and statistics can be viewed under Monitor→PTP.

PTP External Clock Mode

One_PPS_Mode	Output
External Enable	False
Adjust Method	Auto
Clock Frequency	1
One PPS Domain	0

Parameter	Description
One_PPS_Mode	This Selection box will allow you to select the One_pps_mode configuration. The following values are possible: **Output**: Enable the 1 pps clock output **Disable**: Disable the 1 pps clock output
External Enable	This Selection box will allow you to configure the External Clock output. The following values are possible: **True**: Enable the external clock output **False**: Disable the external clock output
Adjust Method	This Selection box will allow you to configure the Frequency adjustment configuration. **LTC**: Select Local Time Counter (LTC) frequency control **Auto**: AUTO Select clock control, based on PTP profile and available HW resources.
Clock Frequency	This will allow to set the Clock Frequency. The possible range of values are **1 - 25000000** (1Hz - 25MHz)
One PPS Domain	This box selects clock domain of chip from which 1pps signal need to be generated. On single clock domain boards, only clock domain zero is available. On boards with multiple clock domains one of three clock domains can be chosen.

PTP timestamping mode

PHY Timestamping Mode	True

This selection box allows the user to disable or enable PHY timestamping on the board.
On whichever ports PHY timestamping is enabled, those ports are affected by this selection.
After disabling PHY timestamping, switch or Mac timestamping is used on those ports.
This option must be saved to start-up config and device must be rebooted for this selection to take effect.

PTP Clock Configuration

Delete	Clock Instance	Clk Domain	VID	Device Type	Profile
Delete	0	0	0	Ord-Bound	No Profile

Parameter	Description

Delete	Check this box and click on 'Save' to delete the clock instance.
Clock Instance	Indicates the instance number of a particular Clock Instance **[0..3]**. Click on the Clock Instance number to edit the Clock details.
Clk Domain	Indicates the clock domain used by the clock.
VID	VLAN Identifier used for tagging the VLAN packets.
Device Type	Indicates the Type of the Clock Instance. There are seven Device Types. **Ord-Bound** - clock's Device Type is Ordinary-Boundary Clock. **P2p Transp** - clock's Device Type is Peer to Peer Transparent Clock. **E2e Transp** - clock's Device Type is End to End Transparent Clock. **Master Only** - clock's Device Type is Master Only. **Slave Only** - clock's Device Type is Slave Only. **AED GM** - clock's Device Type is Grandmaster only (AED profile specific). **Internal** - clock's Device Type is Internal i.e. this device is used to synchronize clock domains on same board.
Profile	Indicates the profile used by the clock.

4.33 MRP (Multiple Registration Protocol)

Multiple Registration Protocol is a generic registration framework that defines the dynamic registration and de-registration of attributes across a Bridged Local Area Network. Such attributes could for example be VLAN identifiers or multicast group MAC addresses. The standard was originally defined by IEEE 802.1ak, and its latest incorporation is in IEEE 802.1Q-2014.

4.33.1 Ports

This page allows you to configure the MRP generic settings for all switch ports.

MRP Overall Port Configuration

Port	Join Timeout	Leave Timeout	LeaveAll Timeout	Periodic Transmission
*	20	60	1000	☐
1	20	60	1000	☐
2	20	60	1000	☐
3	20	60	1000	☐
4	20	60	1000	☐
5	20	60	1000	☐
6	20	60	1000	☐
7	20	60	1000	☐
8	20	60	1000	☐
9	20	60	1000	☐
10	20	60	1000	☐
11	20	60	1000	☐
12	20	60	1000	☐

Parameter	Description
Port	The port number for which the following configuration applies.
Join Timeout	Controls the timeout of the Join Timer for all MRP Applications on this switch port. This value is restricted to 1-20 centiseconds.
Leave Timeout	Controls the timeout of the Leave Timer for all MRP Applications on this switch port. This value is restricted to 60- 300 centiseconds.
LeaveAll Timeout	Controls the timeout of the LeaveAll Timer for all MRP Applications on this switch port. This value is restricted to 1000- 5000 centiseconds.

Periodic Transmission	Enable or disable the Periodic Transmission feature for all MRP Applications on this switch port.

4.33.2 MVRP

Multiple Vlan Registration Protocol (MVRP) is a protocol that defines the dynamic registration and de-registration of VLAN identifiers across a Bridged Local Area Network. It is using the MRP framework to define its operation and therefore it is also called an MRP Application. The standard was originally defined by IEEE 802.1ak, and its latest incorporation is in IEEE 802.1Q-2014. This page allows you to configure the MVRP global and per port settings altogether.
MVRP statistics can be viewed at Monitor→MVRP.

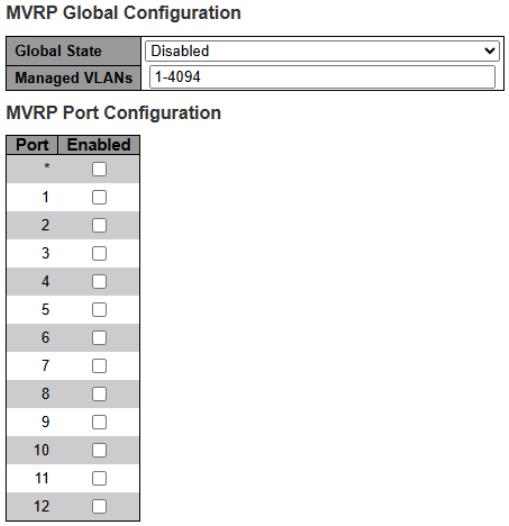

Parameter	Description
Global State	Enable or disable the MVRP protocol globally. This will enable or disable the protocol globally and at the same time on the switch ports that are MVRP enabled.
Managed VLANs	This field shows the managed VLANs, i.e. the VLANs that MVRP will operate upon. By default, only VLANs 1- 4094 are managed, i.e. the entire range as defined in IEEE802.1Q-2014 for MVRP. However this range can be limited by using a list syntax where the individual elements are separated by commas. Ranges are specified with a dash separating the lower and upper bound. The following example will create VLANs 1, 10, 11, 12, 13, 200, and 300: 1,10-13,200,300. Spaces are allowed in between the delimiters.
Port	The port number for which the following configuration applies.
Enabled	Enable or disable the MVRP protocol on this switch port. This will enable or disable the protocol on the switch port given that MVRP is also globally enabled.

4.34 GVRP

GARP VLAN Registration Protocol (GVRP) is a standards-based protocol that automates the creation and management of VLANs across a network of switches and is specified in IEEE 802.1Q-2005, clause 11. It allows a switch to dynamically propagate VLAN registration information to other bridges, ensuring that frames are only forwarded to ports where a destination device is active.

4.34.1 Global config

This page allows you to configure the global GVRP configuration settings that are commonly applied to all GVRP enabled ports.

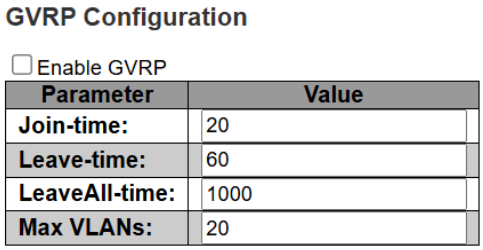

Parameter	Description
Enable GVRP	The GVRP feature is globally enabled by setting the check mark in the checkbox named Enable GVRP and pressing the Save button.
Join-time	Join-time is a value in the range of **1-20cs**, i.e. in units of one hundredth of a second. The default value is 20cs.
Leave-time	Leave-time is a value in the range of **60-300cs**, i.e. in units of one hundredth of a second. The default is 60cs.
LeaveAll-time	LeaveAll-time is a value in the range of **1000-5000cs**, i.e. in units of one hundredth of a second. The default is 1000cs.
Max VLANs	When GVRP is enabled, a maximum number of VLANs supported by GVRP is specified. By default, this number is 20. This number can only be changed when GVRP is turned off.

4.34.2 Port config

This Page allows to enable or disable a port for GVRP operation. The configuration can be performed either before or after GVRP is configured globally – the protocol operation will be the same.

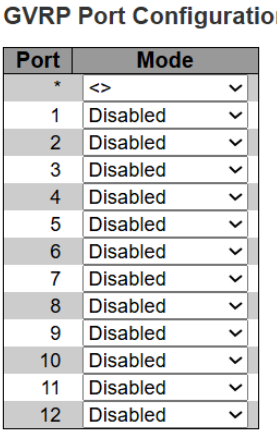

Parameter	Description
Port	The logical port that is to be configured.
Mode	Mode can be either **Disabled** or **GVRP enabled**. These values turn the GVRP feature off or on respectively for the port in question.

4.35 sFlow

sFlow is an industry standard technology for monitoring switched networks through random sampling of packets on switch ports and time-based sampling of port counters. The sampled packets and counters (referred to as flow samples and counter samples, respectively) are sent as sFlow UDP datagrams to a central network traffic monitoring server. This central server is called an sFlow receiver or sFlow collector. Additional information can be found at http://sflow.org.
sFlow status can be viewed at Monitor→sFlow.

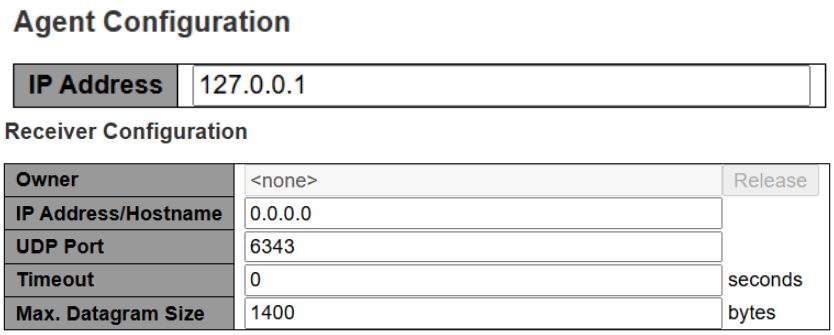

Parameter	Description
IP Address	The IP address used as Agent IP address in sFlow datagrams. It serves as a unique key that will identify this agent over extended periods of time. Both IPv4 and IPv6 addresses are supported.
Owner	Basically, sFlow can be configured in two ways: Through local management using the Web or CLI interface or through SNMP. This read-only field shows the owner of the current sFlow configuration and assumes values as follows: If sFlow is currently unconfigured/unclaimed, Owner contains **<none>**. If sFlow is currently configured through Web or CLI, Owner contains **<Configured through local management>**. If sFlow is currently configured through SNMP, Owner contains a string identifying the sFlow receiver. If sFlow is configured through SNMP, all controls - except for the Release-button - are disabled to avoid inadvertent reconfiguration. The **Release** button allows for releasing the current owner and disable sFlow sampling. The button is disabled if sFlow is currently unclaimed. If configured through SNMP, the release must be confirmed (a confirmation request will appear).
IP Address/Hostname	The IP address or hostname of the sFlow receiver. Both IPv4 and IPv6 addresses are supported.
UDP Port	The UDP port on which the sFlow receiver listens to sFlow datagrams. If set to 0 (zero), the default port (6343) is used.
Timeout	The number of seconds remaining before sampling stops and the current sFlow owner is released. While active, the current time left can be updated with a click on the Refresh-button. If locally managed, the timeout can be changed on the fly without affecting any other settings. Valid range is **0 to 2147483647** seconds.
Max. Datagram Size	The maximum number of data bytes that can be sent in a single sample datagram. This should be set to a value that avoids fragmentation of the sFlow datagrams. Valid range is **200 to 1468** bytes with default being 1400 bytes.

Port Configuration

Port	Flow Sampler			Counter Poller	
	Enabled	Sampling Rate	Max. Header	Enabled	Interval
*	☐	0	128	☐	0
1	☐	0	128	☐	0
2	☐	0	128	☐	0
3	☐	0	128	☐	0
4	☐	0	128	☐	0
5	☐	0	128	☐	0
6	☐	0	128	☐	0
7	☐	0	128	☐	0
8	☐	0	128	☐	0
9	☐	0	128	☐	0
10	☐	0	128	☐	0
11	☐	0	128	☐	0
12	☐	0	128	☐	0

Parameter	Description
Port	The port number for which the configuration below applies.
Flow Sampler Enabled	Enables/disables flow sampling on this port.
Flow Sampler Sampling Rate	The statistical sampling rate for packet sampling. Set to N to sample on average 1/Nth of the packets transmitted/received on the port. Not all sampling rates are achievable. If an unsupported sampling rate is requested, the switch will automatically adjust it to the closest achievable. This will be reported back in this field. Valid range is **`1 to 32767`**.
Flow Sampler Max. Header	The maximum number of bytes that should be copied from a sampled packet to the sFlow datagram. Valid range is **`14 to 200`** bytes with default being 128 bytes. To have room for any frame, the maximum datagram size should be roughly 100 bytes larger than the maximum header size. If the maximum datagram size does not take into account the maximum header size, samples may be dropped.
Counter Poller Enabled	Enables/disables counter polling on this port.
Counter Poller Interval	With counter polling enabled, this specifies the interval - in seconds - between counter poller samples. Valid range is **`1 to 3600`** seconds.

4.36 DDMI

Digital Diagnostics Monitoring Interface (DDMI) provides an enhanced digital diagnostic monitoring interface for optical transceivers which allows real-time access to device operating parameters, if the used SFP modules support DDMI.

DDMI can be viewed at Monitor→DDMI.

DDMI Configuration

Mode	Disabled ▾

Indicates the DDMI mode operation. Possible modes are:

`Enabled`: Enable DDMI mode operation.

`Disabled`: Disable DDMI mode operation.

4.37 UDLD

Uni Directional Link Detection (UDLD) is a Layer 2 protocol that detects and mitigates unidirectional link faults on physical connections. It periodically sends detection messages on enabled ports to verify bidirectional communication with the peer device.

UDLD status can be viewed at Monitor→UDLD.

UDLD Port Configuration

Port	UDLD mode	Message Interval
*	<>	7
1	Disable	7
2	Disable	7
3	Disable	7
4	Disable	7
5	Disable	7
6	Disable	7
7	Disable	7
8	Disable	7
9	Disable	7
10	Disable	7
11	Disable	7
12	Disable	7

Parameter	Description
Port	Port number of the switch.
UDLD Mode	Configures the UDLD mode on a port. Valid values are **`Disable`**, **`Normal`** and **`Aggressive`**. Default mode is Disable. **`Disable`**: UDLD functionality doesn't exist on port. **`Normal`**: If the link state of the port was determined to be unidirectional, it will not affect the port state. **`Aggressive`**: Unidirectional detected ports will get shutdown. To bring back the ports up, need to disable UDLD on that port.
Message Interval	Configures the period of time between UDLD probe messages on ports that are in the advertisement phase and are determined to be bidirectional. The range is from **`7 to 90`** seconds (Default value is 7 seconds)

4.38 Router

This page provides basic routing-related security and control mechanisms. The device allows authentication of routing protocol messages and filtering of routed traffic.

These functions help ensure that routing information is exchanged only between trusted devices and that network traffic is processed according to defined access policies.

4.38.1 Key-Chain

This is router key chain name table. The maximum number of the router key-chain name entries is 64.

Router Key-Chain Configuration

Delete	Key Chain Name	Key ID
	*	*
No entry exists		

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
Key Chain Name	The name of the key-chain entry. The valid name string length is from 1 to 31 and allows all printable characters excluding space character.
Key ID	Click the icon to edit the key.

4.38.2 Key-Chain Key ID

This is router key chain key ID configuration table.

Router Key-Chain Key IDs Configuration

VLAN ID All

Delete	Key Chain Name	Key ID	Change Key String
	*	*	*
Delete		1	

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
Key Chain Name	The name of the key-chain entry. The valid name string length is from 1 to 31 and allows all printable characters excluding space character.
Key ID	The key chain key ID. The valid range is from **1 to 255**.
Change Key String	The key string. It is used to change the key string (fill with plain text). The valid string length is from 1 to 63.

4.38.3 Access-List

This is router access-list configuration table. The maximum number of the router access-list entries is 130.

Router Access-List Configuration

Delete	Name	Mode	Network Address	Mask Length
	*	*	*	*
Delete		Permit	0.0.0.0	0

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
Name	The name of the access-list entry. The valid name string length is from 1 to 31 and allows all printable characters excluding space character.
Mode	The access right mode of the access-list entry. **Permit**: Permit the access right. **Deny**: Deny the access right.

Network Address	The IPv4 address of the access-list entry.
Mask Length	The network prefix size of the access-list entry.

4.39 OSPF

Open Shortest Path First (OSPF) is a link-state routing protocol used to find the best path for data packets as they move through a set of connected networks. It builds a complete topology map of the network by exchanging Link-State Advertisements (LSAs) between routers and uses the Shortest Path First (SPF) algorithm to calculate the most efficient, loop-free routes.

4.39.1 Configuration

This is OSPF router status table. It is used to provide the OSPF router status information.
OSPF status can be viewed at Monitor→OSPF.

OSPF Global Configuration — Clear OSPF Process

OSPF Router Mode			Enable
Router ID			Auto 192.168.203.160 Specific 0.0.0.1
Default Passive Mode			False
Default Metric			Auto Specific 0
Redistribute	**Static**	**Metric Type**	None
		Metric Value	Auto Specific 0
	Connected	**Metric Type**	None
		Metric Value	Auto Specific 0
	RIP	**Metric Type**	None
		Metric Value	Auto Specific 0
Stub Router	**On Startup**	**Mode**	Disable
		Interval	5
	On Shutdown	**Mode**	Disable
		Interval	5
	Administrative Mode		Disable
Default Route Redistribution		**Metric Type**	None
		Metric Value	Auto Specific 0
		Always	Disable
Administrative Distance			110

Parameter	Description
OSPF Router Mode	Enable/Disable the OSPF router mode.
Router ID	The OSPF Router ID in IPv4 address format (A.B.C.D). When the router's OSPF Router ID is changed, if there is one or more fully adjacent neighbors in current OSPF area, the new router ID will take effect after restart OSPF process. Notice that the router ID should be unique in the Autonomous System and value '0.0.0.0' is invalid since it is reserved for the default algorithm. **Auto**: The default algorithm will choose the largest IP address assigned to the router. **Specific**: User specified router ID. The allowed range is from 0.0.0.1 to 255.255.255.254.
Default Passive Mode	Configure all interfaces as passive-interface by default. When an interface is configured as a passive-interface, the OSPF routing updates sending is suppressed, therefore the interface does not establish adjacencies (No OSPF Hellos). The subnet of all interfaces (both passive and active) is advertised by the OSPF router.

Default Metric	User specified default metric value for the OSPF routing protocol. The field is significant only when the argument 'IsSpecificDefMetric' is TRUE. **Auto**: The default metric is calculated automatically based on the routing protocols. **Specific**: User specified default metric. The allowed range is **0 to 16777214**.
Static Redistribute Metric Type	The OSPF redistributed metric type for the static routes. **None**: The static routes are not redistributed. **External Type 1**: External Type 1 of the static routes. **External Type 2**: External Type 2 of the static routes.
Static Redistribute Metric Value	User specified metric value for the static routes. The field is significant only when the argument 'StaticRedistIsSpecificMetric' is TRUE. The allowed range is **0 to 16777214**. **Auto**: The redistributed metric is the same as the original metric value. **Specific**: User specified metric for the static routes.
Connected Redistribute Metric Type	The OSPF redistributed metric type for the connected interfaces. **None**: The connected interfaces are not redistributed. **External Type 1**: External Type 1 of the connected interfaces routes. **External Type 2**: External Type 2 of the connected interfaces routes.
Connected Redistribute Metric Value	User specified metric value for the connected interfaces. The field is significant only when the argument 'ConnectedRedistIsSpecificMetric' is TRUE. The allowed range is **0 to 16777214**. **Auto**: The redistributed metric is the same as the original metric value. **Specific**: User specified metric for the connected routes.
RIP Redistribute Metric Type	The OSPF redistributed metric type for the RIP routes. The field is significant only when the RIP protocol is supported on the device. **None**: The RIP routes are not redistributed. **External Type 1**: External Type 1 of the RIP routes. **External Type 2**: External Type 2 of the RIP routes.
RIP Redistribute Metric Value	User specified metric value for the RIP routes. The field is significant only when the RIP protocol is supported on the device and argument 'RipRedistIsSpecificMetric' is TRUE.. The allowed range is **0 to 16777214**. **Auto**: The redistributed metric is the same as the original metric value. **Specific**: User specified metric for the RIP routes.
Stub router during startup period	Configures OSPF to advertise a maximum metric during startup for a configured period of time.
Stub router on startup interval time	User specified time interval (seconds) to advertise itself as stub area. The field is significant only when the on-startup mode is enabled.. The allowed range is **5 to 86400** seconds.
Stub router during shutdown period	Configures OSPF to advertise a maximum metric during shutdown for a configured period of time. The device advertises a maximum metric when the OSPF router mode is disabled and notice that the mechanism also works when the device reboots but not for the 'reload default' case.
Stub router on shutdown interval time	User specified time interval (seconds) to wait till shutdown completed. The field is significant only when the on-shutdown mode is enabled.. The allowed range is **5 to 100** seconds.
Stub router administrative mode	Configures OSPF stub router mode administratively applied, for an indefinite period.
Default Route Redistribution Metric Type	The OSPF redistributed metric type for a default route. **None**: The default route are not redistributed. **External Type 1**: External Type 1 of the default route. **External Type 2**: External Type 2 of the default route.
Default Route Redistribution Metric Value	User specified metric value for a default route. The field is significant only when the argument 'DefaultRouteRedistIsSpecificMetric' is TRUE. The allowed range is **0 to 16777214**. **Auto**: The redistributed metric is the same as the original metric value. **Specific**: User specified metric for the default route.
Default Route Redistribution Always	Specifies to always advertise a default route into all external-routing capable areas. Otherwise, the router only advertises the default route when the advertising router already has a default route.
Administrative Distance	The OSPF administrative distance.

4.39.2 Network Area

This is OSPF area configuration table. It is used to specify the OSPF enabled interface(s). When OSPF is enabled on the specific interface(s), the router can provide the network information to the other OSPF routers via those interfaces.
OSPF status can be viewed at Monitor→OSPF.

OSPF Network Area Configuration

Delete	Network Address	Mask Length	Area ID
	*	*	*
Delete	0.0.0.0	24	0.0.0.0

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
Network Address	IPv4 network address.
Mask Length	IPv4 network mask length.
Area ID	The OSPF area ID.

4.39.3 Passive Interface

This is OSPF router interface configuration table.
OSPF status can be viewed at Monitor→OSPF.

OSPF Passive Interface Configuration

Interface	Passive Interface
*	☐
VLAN 1	☐

Parameter	Description
Interface	Interface identification.
Passive Interface	Enable the interface as OSPF passive interface.

4.39.4 Stub Area

This is OSPF stub area configuration table. The configuration is used to reduce the link-state database size and therefore the memory and CPU requirement by forbidding some LSAs.
OSPF status can be viewed at Monitor→OSPF.

OSPF Area Stub Configuration

Delete	Area ID	Stub Type	No Summary	Translator Role
	*			
Delete	0.0.0.0	Stub Area	☐	Candidate

Parameter	Description

Delete	Check to delete the entry. It will be deleted during the next save.
Area ID	The OSPF area ID.
Stub Type	The OSPF stub configured type. **Stub Area**: Configure the area as stub area. **NSSA**: Configure the area as not-so-stubby area (NSSA).
No Summary	The value is true to configure the inter-area routes do not inject into this stub area.
Translator Role	The OSPF NSSA translator role. **Candidate**: this NSSA-ABR router will participate in the translator election. **Never**: this NSSA-ABR router never translates. **Always**: this NSSA-ABR router always translates.

4.39.5 Area Authentication

This is OSPF area authentication configuration table. It is used to apply the authentication to all the interfaces belonging to the area.
OSPF status can be viewed at Monitor→OSPF.

OSPF Area Authentication Configuration

Delete	Area ID	Auth. Type
	*	
Delete	0.0.0.0	Simple Password

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
Area ID	The OSPF area ID.
Auth. Type	The authentication type on an area is applied to all the interfaces belong to that area. The authentication type on an IP interface or a virtual link overrides the authentication type on an area and is useful if different interfaces in the same area use different authentication types. Specify the authentication type. **Simple Password**: Simple password authentication. **Message Digest**: MD5 digest authentication.

4.39.6 Area Range

This is OSPF area range configuration table. It is used to summarize the intra area paths from a specific address range in one summary-LSA (Type-3) and advertised to other areas or configure the address range status as 'DoNotAdvertise' which the summary-LSA (Type-3) is suppressed. The area range configuration is used for Area Border Routers (ABRs) and only router-LSAs (Type-1) and network-LSAs (Type-2) can be summarized. The AS-external-LSAs (Type-5) cannot be summarized because the scope is OSPF autonomous system (AS). The AS-external-LSAs (Type-7) cannot be summarized because the feature is not supported yet.
OSPF status can be viewed at Monitor→OSPF.

OSPF Area Range Configuration

Delete	Area ID	Network Address	Mask Length	Advertise	Cost	
	*	*	*			
Delete	0.0.0.0	0.0.0.0	24	☑	Auto	0

Parameter	Description

Delete	Check to delete the entry. It will be deleted during the next save.
Area ID	The OSPF area ID.
Network Address	IPv4 network address.
Mask Length	IPv4 network mask length.
Advertised	When the value is true, it summarizes intra area paths from the address range in one summary-LSA (Type-3) and advertised to other areas. Otherwise, the intra area paths from the address range are not advertised to other areas.
Auto/Specific	When 'Auto' is selected, the cost value is set to 0 automatically and isn't allowed to be configured.
Cost	User specified cost (or metric) for this summary route. It is allowed to be configured only when 'Specific' is selected. The allowed range is **0 to 16777215** and the default setting is 'auto cost' mode.

4.39.7 Interfaces

This is interface configuration parameter table.
OSPF status can be viewed at Monitor→OSPF.

OSPF Interface Configuration

Interface	Priority	Cost		Ignore MTU	FastHelloPackets	Interval			Auth. Type	Change Simple Password		MD Key
						Hello	Dead	Retransmit				
*	1	<>	1	☐	☐	2	10	40	5	<>	*	*	*
VLAN 1	1	Auto	1	☐	☐	2	10	40	5	Area Configuration	☐		◎

Parameter	Description
Interface	Interface identification.
Priority	User specified router priority for the interface. The allowed range is **0 to 255** and the default value is 1.
Cost	User specified cost for this interface. It's link state metric for the interface. The field is significant only when 'IsSpecificCost' is TRUE. The allowed range is **1 to 65535** and the default setting is 'auto cost' mode.
Ignore MTU	When false (default), the interface MTU of received OSPF database description (DBD) packets will be checked against our own MTU, and if the remote MTU is greater than our own, the DBD will be ignored. When true, the interface MTU of the received OSPF DBD packets will be ignored.
FastHelloPackets	How many Hello packets will be sent per second. The allowed range is **1 to 10** and the default setting is disabled.
Hello Interval	How many Hello packets will be sent per second. The allowed range is **1 to 65535** and the default value is 10 (seconds).
Dead Interval	The time interval (in seconds) between hello packets. The allowed range is **1 to 65535** and the default value is 40 (seconds).
Retransmit Interval	The time interval (in seconds) between link-state advertisement (LSA) retransmissions for adjacencies. The allowed range is **3 to 65535** and the default value is 5 (seconds).
Auth. Type	The authentication type. **Simple Password**: It's using a plain text authentication. A password must be configured, but the password can be read by sniffer the packets. **Message Digest**: It's message-digest algorithm 5 (MD5) authentication. Keying material must also be configured. This is the most secure method. **Null Authentication**: No authentication. **Area Configuration**: Refer to Area authentication setting.
Change Simple Password	It is used to change the simple password (fill with plain text). The allowed input length is **1 to 8**.
MD Key	Click the icon to edit the message digest key for the entry.

OSPF Interface Message Digest Configuration

Interface: VLAN 1

Delete	Interface	MD Key ID	Password
	*	*	*
Delete	VLAN 1	1	

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
Interface	Interface identification.
MD Key ID	The key ID for message digest authentication. The allowed range is **1 to 255**.
Password	The message digest key. The allowed input length is **1 to 16**.

4.39.8 Virtual Link

This is OSPF virtual link configuration table. The virtual link is established between 2 ABRs to overcome that all the areas have to be connected directly to the backbone area.
OSPF status can be viewed at Monitor→OSPF.

OSPF Virtual Link Configuration

Delete	Area ID	Router ID	Interval			Auth. Type	Change Simple Password	MD Key
			Hello	Dead	Retransmit			
	*	*					*	*
Delete	0.0.0.0	0.0.0.0	10	40	5	Area Configuration		

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
Area ID	OSPF Area ID.
Router ID	OSPF router ID.
Hello Interval	The time interval (in seconds) between hello packets. The allowed range is **1 to 65535** and the default value is 10 (seconds).
Dead Interval	The time interval (in seconds) between hello packets. The allowed range is **1 to 65535** and the default value is 40 (seconds).
Retransmit Interval	The time interval (in seconds) between link-state advertisement (LSA) retransmissions for adjacencies. The allowed range is **3 to 65535** and the default value is 5 (seconds).
Auth. Type	The authentication type on an area. **Simple Password**: It's using a plain text authentication. A password must be configured, but the password can be read by sniffer the packets. **Message Digest**: It's message-digest algorithm 5 (MD5) authentication. Keying material must also be configured. This is the most secure method. **Null Authentication**: No authentication. **Area Configuration**: Refer to Area authentication setting.
Change Simple Password	It is used to change the simple password (fill with plain text). The allowed input length is **1 to 8**.
MD Key	Click the icon to edit the message digest key for the entry.

OSPF Virtual Link Message Digest Configuration

Area ID: 192.168.0.20, Router ID: 192.168.0.10

Delete	Area ID	Router ID	MD Key ID	Password
	*	*	*	*
Delete	192.168.0.20	192.168.0.10	1	

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
Area ID	OSPF Area ID.
Router ID	OSPF router ID.
MD Key ID	The key ID for message digest authentication. The allowed range is **1 to 255**.
Password	The message digest key. The allowed input length is **1 to 16**.

4.40 OSPFv3

OSPFv3 is the evolution of the OSPF protocol specifically redesigned to support IPv6, though it can also carry IPv4 traffic through the use of multi-topology extensions. Unlike its predecessor, OSPFv3 runs on a per-link basis rather than a per-subnet basis, allowing multiple address prefixes to exist on a single interface.

4.40.1 Configuration

This is OSPF6 router configuration table. It is a general group to configure the OSPF6 common router parameters.
OSPFv3 status information and statistics can be viewed at Monitor→OSPFv3.

OSPF6 Global Configuration

OSPF6 Router Mode		Enable
Router ID		Auto 192.168.203.160 Specific 0.0.0.1
Redistribute	Static	Disable
	Connected	Disable
Administrative Distance		110

Parameter	Description
OSPF6 Router Mode	Enable/Disable the OSPF6 router mode.
Router ID	The OSPF6 Router ID in IPv4 address format(A.B.C.D). When the router's OSPF6 Router ID is changed, if there is one or more fully adjacent neighbors in current OSPF6 area, the new router ID will take effect after restart OSPF6 process. Notice that the router ID should be unique in the Autonomous System and value '0.0.0.0' is invalid since it is reserved for the default algorithm. **Auto**: The default algorithm will choose the largest IP address assigned to the router. **Specific**: User specified router ID. The allowed range is from 0.0.0.1 to 255.255.255.254.
Static Redistribute	The OSPF redistribute setting for the static routes. **Enable**: The static routes are redistributed. **Disable**: The static routes are not redistributed

Connected Redistribute	The OSPF redistribute enabled for connected route or not. **Enable**: The connected interfaces are redistributed. **Disable**: The connected interfaces are not redistributed.
Administrative Distance	The OSPF6 administrative distance.

4.40.2 Interface Area

This is OSPF6 router interface configuration table.
The area status can be viewed at Monitor→OSPFv3→Area.

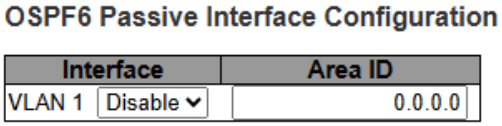

Parameter	Description
Interface	Interface identification.
Interface Area ID	The OSPF6 interface Area ID. Only valid if 'is_specific_id' is true

4.40.3 Stub Area

This is OSPF6 stub area configuration table. The configuration is used to reduce the link-state database size and therefore the memory and CPU requirement by forbidding some LSAs.
The area status can be viewed at Monitor → OSPFv3 → Area.

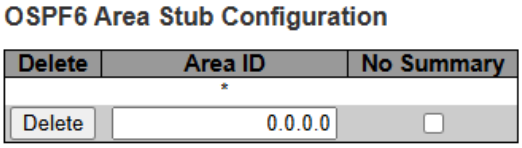

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
Area ID	The OSPF6 area ID.
No Summary	The value is true to configure the inter-area routes do not inject into this stub area.

4.40.4 Area Range

This is OSPF6 area range configuration table. It is used to summarize the intra area paths from a specific address range in one summary-LSA (Type-0x2003) and advertised to other areas or configure the address range status as 'DoNotAdvertise' which the summary-LSA(Type-0x2003) is suppressed. The area range configuration is used for Area Border Routers (ABRs), and only router-LSAs (Type-0x2001) and network-LSAs (Type-0x2002) can be summarized. The AS-external-LSAs (Type-0x4005) cannot be summarized because the scope is OSPF6 autonomous system (AS). The AS-external-LSAs (Type-0x4007) cannot be summarized.
The area status can be viewed at Monitor→OSPFv3→Area.

OSPF6 Area Range Configuration

Delete	Area ID	Network Address	Mask Length	Advertise	Cost
	*	*	*		
Delete	0.0.0.0	0::0	128	☑	Auto 0

Add New Entry

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
Area ID	The OSPF6 area ID.
Network Address	IPv6 network address.
Mask Length	IPv6 network mask length.
Advertised	When the value is true, it summarizes intra area paths from the address range in one Inter-Area Prefix LSA (Type-0x2003) and advertised to other areas. Otherwise, the intra area paths from the address range are not advertised to other areas.
Auto/Specific	When 'Auto' is selected, the cost value is set to 0 automatically and isn't allowed to be configured.
Cost	User specified cost (or metric) for this summary route. It is allowed to be configured only when 'Specific' is selected. The allowed range is **0 to 16777215** and the default setting is 'auto cost' mode.

4.40.5 Interfaces

This is interface configuration parameter table.
OSPFv3 Interface status can be viewed at Monitor→OSPFv3→Interface.

OSPF6 Interface Configuration

Interface	Priority	Passive Interface	Cost		Ignore MTU	Interval		
						Hello	Dead	Retransmit
*	1	☐	<>	1	☐	10	40	5
VLAN 1	1	☐	Auto	1	☐	10	40	5

Parameter	Description
Interface	Interface identification.
Priority	User specified router priority for the interface. The allowed range is **0 to 255** and the default value is 1.
Passive Interface	Indicates whether the interface is passive or not
Cost	User specified cost for this interface. It's link state metric for the interface. The field is significant only when 'IsSpecificCost' is TRUE. The allowed range is **1 to 65535** and the default setting is 'auto cost' mode.
Ignore MTU	When false (default), the interface MTU of received OSPFv3 database description (DBD) packets will be checked against our own MTU, and if the remote MTU differs from our own, the DBD will be ignored. When true, the interface MTU of the received OSPFv3 DBD packets will be ignored.
Hello Interval	How many Hello packets will be sent per second. The allowed range is **1 to 65535** and the default value is 10 (seconds).
Dead Interval	The time interval (in seconds) between hello packets. The allowed range is **1 to 65535** and the default value is 40 (seconds).
Retransmit Interval	The time interval (in seconds) between link-state advertisement (LSA) retransmissions for adjacencies. The allowed range is **3 to 65535** and the default value is 5 (seconds).

4.41 RIP

Routing Information Protocol (RIP) is a distance-vector routing protocol that uses hop count as its primary metric to determine the best path between source and destination networks. It periodically broadcasts or multicasts its entire routing table to neighboring devices, which then update their own tables based on the shortest path discovered. To prevent routing loops, RIP limits the maximum path length to 15 hops, considering any destination 16 hops away as unreachable.

4.41.1 Configuration

This is RIP router configuration table. It is a general group to configure the RIP common router parameters. RIP status and additional information can be viewed at Monitor→RIP.

RIP Global Configuration

RIP Router Mode			Disable
Version			Default
Timers	Update		30
	Invalid		180
	Garbage-Collection		120
Redistribute	Static	Mode	Disable
		Metric Value	Auto / Specific 1
	Connected	Mode	Disable
		Metric Value	Auto / Specific 1
	OSPF	Mode	Disable
		Metric Value	Auto / Specific 1
	Default Metric Value		1
	Default Route		Disable
Default Passive Mode			Disable
Administrative Distance			120

Parameter	Description
RIP Router Mode	Enable/Disable the RIP router mode. **Enable**: Enable the the RIP router mode. **Disable**: Disable the the RIP router mode.
Version	RIP version support. **Default**: Base on the default version process. The router sends RIPv2 and accepts both RIPv1 and RIPv2. When the router receives either version of REQUESTS or triggered updates packets, it replies with the appropriate version. **Version 1**: Receive/Send RIPv1 only. **Version 2**: Receive/Send RIPv2 only.
Update Timer	The timer interval (in seconds) between the router sends the complete routing table to all neighboring RIP routers. The allowed range is **5 to 2147483**.
Invalid Timer	The invalid timer is the number of seconds after which a route will be marked invalid. The allowed range is **5 to 2147483**.
Garbage Collection Timer	The garbage collection timer is the number of seconds after which a route will be deleted. The allowed range is **5 to 2147483**.
Static Redistribute Mode	Indicate if the router redistribute the static routes into the RIP domain or not. **Enable**: Enable static routes redistribution. **Disable**: Disable static routes redistribution.
Static Redistribute Metric Value	User specified metric value for the static routes. The field is significant only when the argument 'StaticRedistIsSpecificMetric' is TRUE. If the specific metric setting is removed while the static redistributed mode is enabled, the router will update the original static redistributed routes with metric value 16 before updates to the new metric value. The allowed range is **1 to 16**.

	Auto: The redistributed metric value is refer to redistributed default metric value. **Specific**: User specified metric for the static routes.
Connected Redistribute Mode	This setting determines whether the router shares its "Directly Connected" routes (subnets physically attached to its interfaces) with other RIP neighbors. **Enable**: Enable connected routes redistribution. **Disable**: Disable connected routes redistribution.
Connected Redistribute Metric Value	User specified metric value for the connected interfaces. The field is significant only when the argument 'ConnectedRedistIsSpecificMetric' is TRUE. If the specific metric setting is removed while the connected redistributed mode is enabled, the router will updates the original connected redistributed routes with metric value 16 before updates to the new metric value. The allowed range is **1 to 16**. **Auto**: The redistributed metric value is refer to redistributed default metric value. **Specific**: User specified metric for the connected routes.
OSPF Redistribute Mode	Indicate if the router redistribute the OSPF routes into the RIP domain or not. The field is significant only when the OSPF protocol is supported on the device. **Enable**: Enable OSPF routes redistribution. **Disable**: Disable OSPF routes redistribution.
OSPF Redistribute Metric Value	User specified metric value for the RIP routes. The field is significant only when the OSPF protocol is supported on the device and argument 'OspfRedistIsSpecificMetric' is TRUE. If the specific metric setting is removed while the OSPF redistributed mode is enabled, the router will updates the original OSPF redistributed routes with metric value 16 before updates to the new metric value The allowed range is **1 to 16**. **Auto**: The redistributed metric value is refer to redistributed default metric value. **Specific**: User specified metric for the OSPF routes.
Redistribute Default Metric Value	The RIP default redistributed metric. It is used when the metric value isn't specified for the redistributed protocol type. The allowed range is **1 to 16**.
Default Passive Mode	The RIP default route redistribution.
Redistribute Default Route	Configure all interfaces as passive-interface by default.
Administrative Distance	The RIP administrative distance. The allowed range is **1 to 255**.

4.41.2 Network

This is RIP network configuration table. It is used to specify the RIP enabled interface(s). When RIP is enabled on the specific interface(s), the router can provide the network information to the other RIP routers via those interfaces. The maximum number of RIP network segment entries is 126.
RIP status and additional information can be viewed at Monitor→RIP.

RIP Network Configuration

Delete	Network Address	Mask Length
	*	*
Delete	0.0.0.0	24

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
Network Address	IPv4 network address.
Mask Length	IPv4 network mask length.

4.41.3 Neighbors

This is RIP neighbor connection table. It is used to configure the RIP router to send RIP updates to specific neighbors using the unicast, broadcast, or network IP address after update timer expiration. The maximum number of RIP neighbor entries is 128.
RIP status and additional information can be viewed at Monitor→RIP.

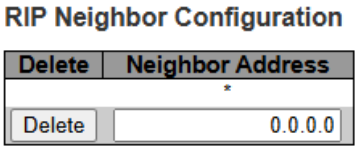

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
Neighbor Address	Ipv4 address encoded as **a.b.c.d**, where a-d is a base-10 human readable integer in the range [0-255] The neighbor address can be an unicast (excluding loopback), broadcast, or network IP address.

4.41.4 Passive Interface

This is RIP router passive interface configuration table.
RIP Interface status can be viewed at Monitor→RIP→Interface.

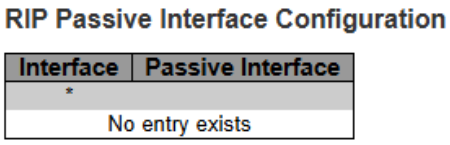

Parameter	Description
Interface	Interface identification.
Passive Interface	Enable the interface as RIP passive-interface.

4.41.5 Interfaces

This is the RIP interface configuration table.
RIP Interface status can be viewed at Monitor→RIP→Interface.

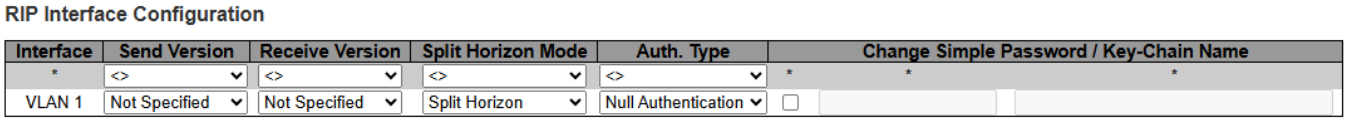

Parameter	Description
Interface	Interface identification.
Send Version	The RIP version for the advertisement transmission on the interface.
Receive Version	The RIP version for the advertisement reception on the interface.

Split Horizon Mode	The split horizon mode. **Split Horizon**: To omit routes learned from one neighbor in updates sent to that neighbor. **Poisoned Reverse**: The neighbor learned routes are included in updates sent to the neighbors but their metrics are set to infinity. **Disabled**: Split horizon is disabled.
Auth. Type	The authentication type. **Simple Password**: It's using a plain text authentication. A password must be configured, but the password can be read by sniffer the packets. **Message Digest**: It's message-digest algorithm 5 (MD5) authentication. Keying material must also be configured. This is the most secure method. **Null Authentication**: No authentication.
Change Simple Password	It is used to change the simple password (fill with plain text). The allowed input length is from 1 to 15 printable characters excluding space character. The null string identifies no simple password is set on the interface. Notice that can not set key chain and simple password at the same time.
Change Key-Chain Name	It is used to change the key chain name used by MD5 authentication. The allowed input length is from 1 to 31 printable characters excluding space character. The null string identifies no key-chain name is set on the interface. Notice that can not set key chain and simple password at the same time.

4.41.6 Offset-List

This is RIP offset-list configuration table. The maximum number of RIP offset-list entries is 130.
RIP status and additional information can be viewed at Monitor→RIP.

RIP Offset-List Configuration

Delete	VLAN ID	Direction	Access List Name	Offset Metric
	*	*	*	*
Delete	0	In		1

Parameter	Description
Delete	Check to delete the entry. It will be deleted during the next save.
VLAN ID	The VLAN interface which the offset list applies to. The range of VLAN ID is from **0 to 4095**. 0 means that the offset list applies to all interfaces.
Direction	The direction to add the offset to routing metric update. **In**: Apply to the inbound direction. **Out**: Apply to the outbound direction.
Access List Name	Access-list name. The valid name string length is from 1 to 31 and allows all printable characters excluding space character.
Offset Metric	The offset to incoming or outgoing routing metric. The allowed range is **0 to 16**.

4.42 RedBox

RedBox stands for Redundancy Box and covers PRP and HSR functionality.
PRP and HSR are specified in IEC62439-3, ED4 (2021) and its CORRIGENDUM 1 (2023).

PRP (Parallel Redundancy Protocol) allows attaching nodes to two separate networks so that connectivity can occur via one or the other network, thereby providing redundancy.

Devices using **HSR** (High-availability Seamless Redundancy) send identical frames in two directions into a so-called HSR ring. Receivers will receive two (or more) copies and discard all but one, thus providing redundancy.

There is one row per possible RedBox instance. A RedBox instance must be created before it is configurable. The buttons in the Action column allow for this.
RedBox status and statistics can be viewed at Monitor→RedBox.

RedBox Configuration

Action	Instance	Enable	Mode	Port A	Port B	Mode U	Net ID	LAN ID	Age Times		
									NodesTable	ProxyNodeTable	Duplicate Discard
Create	1	☐	PRP-SAN			☐	1	A	60	60	40
Create	2	☐	PRP-SAN			☐	1	A	60	60	40
Create	3	☐	PRP-SAN			☐	1	A	60	60	40
Create	4	☐	PRP-SAN			☐	1	A	60	60	40
Create	5	☐	PRP-SAN			☐	1	A	60	60	40

Supervision Frames						
VLAN ID	PCP	DMAC LSByte	Interval	PRP-to-HSR	HSR-to-PRP	Operational State
0	7	0x00	2	☑	☑	
0	7	0x00	2	☑	☑	
0	7	0x00	2	☑	☑	
0	7	0x00	2	☑	☑	
0	7	0x00	2	☑	☑	

Parameter	Description
Action	To create a given RedBox instance, press the Create button. Already created RedBox instances can be deleted with a click on the Delete button. Notice that the click itself doesn't change the underlying configuration. Only when the Save button is clicked, will the configuration take effect.
Instance	Identifies the RedBox instance number.
Enable	Enables or disables a given RedBox instance. If enabling, a click on the Save button causes software to perform additional checks before actually enabling the RedBox instance in hardware. If a check fails, the user will get notified. If disabling, the RedBox instance is removed entirely from hardware.
Mode	A RedBox may run in one of four different modes: **PRP-SAN**: The two LRE ports (Port A and Port B) are connected to the PRP network's LAN A and LAN B, respectively. Port C corresponds to the remaining ports in the switch core and makes up the 'gateway' to the SAN network. Port C is also known as the interlink port. **HSR-SAN**: The two LRE ports (Port A and Port B) connect to the HSR ring and the switch core (Port C) connects to the SAN network. **HSR-PRP**: The two LRE ports connect to the HSR ring managed by this RedBox instance and Port C connects to a PRP network - either LAN A or LAN B. **HSR-HSR**: The two LRE ports connect to the HSR ring managed by this RedBox instance and Port C connects to another RedBox' Port C. You may think of the connection between the two RedBoxes as a wire. The other RedBox may or may not be located on the same device as this RedBox.
Port A	Select the physical interface that constitutes LRE Port A. A given physical port can serve only one specific RedBox instance. The drop-down box lists only those interfaces that the corresponding RedBox supports. It is possible to 'serialize' two or more RedBoxes by utilizing a special port called 'Neighbor'. With this interface, an internal connection is made to connect one RedBox to another, freeing up two physical ports on the device for other purposes. RedBox instance 1's Port A does not have a Neighbor option, because there is no RedBox 'to the left' of RedBox #1. Likewise, RedBox instance 5's Port B does not have a Neighbor option, because there is no RedBox 'to the right' of the last RedBox instance.
Port B	Select the physical interface that constitutes LRE Port B. See description of Port A for details.
Mode U	When checked, unicast frames with a DMAC in the ProxyNodeTable that arrive on an LRE port are - besides being forwarded to the interlink - also forwarded to the other LRE port. Not available in PRP-SAN mode.
Net ID	The Net ID (called 'NetId' in IEC-62439-3) is a number from **1 to 7** and is used in HSR-PRP and HSR-HSR modes, only. The Net ID is along with the LAN ID ('LanId' in IEC-62439-3) used to form the 4-bit Path ID ('PathId' in IEC-62439-3), which is used in frames transmitted to the RedBox' HSR ring.

	In the other direction, it is used to prevent frames originated by this RedBox to return to the interlink and the PRP network (in HSR-PRP mode) or accompanying HSR RedBox (in HSR-HSR mode), by a method known as Net ID filtering: Frames arriving on the LRE ports carrying the same Net ID as the RedBox is configured with will be filtered towards Port C.
LAN ID	The LAN ID (called 'LanId' in IEC-62439-3) is used in HSR-PRP mode, only. It tells the RedBox whether the PRP-side of the RedBox is connected to LAN A or LAN B. In this way, the RedBox can convert the HSR tag in frames arriving on LRE ports to a correct RCT before it is sent to the PRP network.
Nodes Table Age Time	The NodesTable keeps track of source MAC addresses (SMACs) of frames received on LRE ports. In PRP-SAN mode, the NodesTable is used to identify a node either as a SAN or a DAN (DANP). The first time a given SMAC is seen on an LRE port, it is marked as a SAN, so the frames coming from the interlink and destined to that MAC address will be sent to one LRE port, only. If the same SMAC is seen on both Port A and Port B - or a supervision frame is received for that MAC address, the entry will be marked as a DAN, so that frames destined to that MAC address and coming from the interlink will have an RCT added and will be sent to both Port A and Port B. Frames arriving on LRE ports with same <SMAC, RCT.SeqNr> tuple are subject to duplicate discarding. In HSR modes (HSR-SAN, HSR-PRP, and HSR-HSR), the NodesTable serves a slightly different purpose. The first time a given SMAC is seen on an LRE port, that MAC address is added to the NodesTable as a DAN (SANs do not exist on HSR rings in Mode H). Frames destined to that MAC address coming from the interlink are always sent to both Port A and Port B with an HSR tag, and frames arriving on both Port A and Port B are subject to duplicate discarding based on the <SMAC, HSR.SeqNr> tuple. The NodesTable Age Time determines the number of seconds a given entry resides in the NodesTable after it was last seen on Port A or Port B. Valid range is **`1 - 65 seconds`** with a default of 60 seconds.
Proxy Node Table Age Time	The ProxyNodeTable keeps track of source MAC addresses (SMACs) of frames received on the interlink port (Port C). In PRP-SAN and HSR-SAN mode, all nodes in the ProxyNodeTable are considered SANs, and frames coming from the LRE ports and leaving the interlink port will have the RCT (PRP-SAN) or HSR tag (HSR-SAN) removed. In HSR-PRP mode, a node in the ProxyNodeTable can either be a SAN or a DAN (DANP). A supervision frame sent by a DANP in the PRP network will cause the entry type to change from a SAN to a DAN. Once it is identified as a DAN, subsequent frames arriving from the DANP must carry an RCT. If not, those frames will be discarded by the RedBox and only forwarded in the PRP network. The ProxyNodeTable Age Time determines the number of seconds a given entry resides in the ProxyNodeTable after it was last seen on the interlink port. Valid range is **`1 - 65 seconds`** with a default of 60 seconds.
Duplicate Discard Age Time	The Duplicate Discard Age Time determines how long the RedBox shall wait for a copy of <SMAC, SeqNr> before the RedBox considers all copies received and forgets the entry. Valid range is **`10 - 10000 milliseconds`** with a default of 40 milliseconds.
Supervision Frames VLAN ID	The RedBox sends - towards the LRE ports - supervision frames on behalf of SANs connected to the switch core side of the RedBox in PRP-SAN, HSR-SAN, and HSR-PRP mode. This field chooses the VLAN ID with which these supervision frames are sent. If using the value 0, the RedBox sends with the interlink port's native VLAN ID (the Port VLAN ID). The RedBox software knows whether to actually tag the supervision frames based on the interlink port's tagging configuration. It also knows whether the interlink port is a member of the VLAN ID, and if not, it refrains from sending supervision frames. Notice: The RedBox hardware only supports recognizing HSR tags behind VLAN tags with a TPID of 0x8100 (C-tag), so if the interlink port is configured to VLAN tag with another TPID, the software sends with that TPID but receiving RedBoxes may not understand this as a VLAN tag, so it will not be able to find the HSR tag behind it, which may cause such frames to be discarded. This goes for all sorts of frames, not only supervision frames. Valid range is **`0 - 4095`** with a default of 0.
Supervision Frames PCP	Whenever the RedBox software determines that a frames should be transmitted VLAN tagged, it also inserts a PCP value into the VLAN tag. This value comes from this configuration. Notice, that even the native VLAN may have to be transmitted tagged, so the PCP value may also matter when the VLAN ID is set to 0. Valid range is **`0 - 7`** with a default of 7.
Supervision Frames DMAC LSByte	This field controls the value used in the least significant byte of the destination MAC address used in supervision frames sent by the RedBox, that is the 'xx' of 01:15:4E:00:01:xx. A network operator can identify certain parts of the network through the use of different values of this field. Upon reception of supervision frames, the RedBox ignores the LSbyte. Valid range is **`0x00 - 0xFF`** with a default of 0x00.

Supervision Frames Interval	The interval with which supervision frames are sent by this RedBox. To avoid bursting of supervision frames, the first frame sent on behalf of a SAN is transmitted after a random number of milliseconds between 0 and this value. Subsequent supervision frames for a given SAN are transmitted this value seconds apart. Valid range is **`1 - 60 seconds`** with a default of 2 seconds.
Supervision Frames PRP-to-HSR	This field is only used in HSR-PRP mode. When checked, supervision frames from the PRP network are software forwarded as HSR supervision frames to the HSR ring. When unchecked, supervision frames from the PRP network are hardware forwarded without modifications to the HSR ring. Especially older HSR-capable equipment that only supports supervision frames of the HSR type require this option to be checked. Except for changing the supervision frame type from PRP-Duplicate-Discard or PRP-Duplicate-Accept to HSR when software forwarding, more checks are done on the supervision frame, such as checks for correct RCT, correct DMAC (except for least significant byte), correct TLV1.MAC, and correct TLV2.MAC. These checks cannot be performed when hardware forwarding the frames. Default is checked.
Supervision Frames HSR-to-PRP	This field is only used in HSR-PRP mode. When checked, supervision frames from the HSR ring are software forwarded as PRP-Duplicate-Discard supervision frames to the PRP network. When unchecked, supervision frames from the HSR ring are hardware forwarded without modifications to the PRP network. Especially older PRP-capable equipment that only supports supervision frames of the PRP type require this option to be checked. Except for changing the supervision frame type from HSR to PRP-Duplicate-Discard when software forwarding, more checks are done on the supervision frame, such as checks for correct HSR tag, correct DMAC (except for least significant byte), correct TLV1.MAC, and correct TLV2.MAC. These checks cannot be performed when hardware forwarding the frames. Default is checked.
Operational State	This shows the current operational state of a RedBox instance. A color indicates the state as follows:  : The instance is not created or not enabled  : The instance is enabled with no configurational warnings : The instance is enabled, but there are configurational warnings  : The instance is enabled, but an internal error has occurred. See console or crashlog for details When yellow, hover the mouse over the image to see a list of configurational warnings. The possible warnings are as follows: • The MTU is too high on at least one of the LRE ports (max is 2000) • The MTU is too high on at least one non-LRE port. Frames larger than 1994 cannot traverse the HSR/PRP network • Interlink port must use C-tags • Interlink port is not member of the supervision frame VLAN ID • The neighbor RedBox is not configured • The neighbor RedBox is not active • The neighbor's port A is not configured as a RedBox neighbor • The neighbor's port B is not configured as a RedBox neighbor • The neighbor's interlink port has coinciding VLAN memberships with this RedBox's interlink port • Interlink port has spanning tree enabled It is highly recommended to have a green color on enabled instances. It is important to notice that not all configuration errors can be detected by the RedBox software.

5 Monitor

The Monitor section of Web GUI contains important information about the status of active features, network conditions, and the switch itself. Most Monitoring features have a corresponding configuration page in the Configuration section.

The displayed data can be refreshed manually by pressing the Refresh button or automatically every 3 seconds by checking the Auto-refresh ☐ box.
When navigating tables, the >> button will use the last entry of the currently displayed table as a basis for the next lookup. When the end is reached the text *No more entries* are shown in the displayed table. Use the |<< button to start over.

5.1 System

This section contains general system status and information about the switch, IP status, and system logs.

5.1.1 Information

Get current information about the switch.

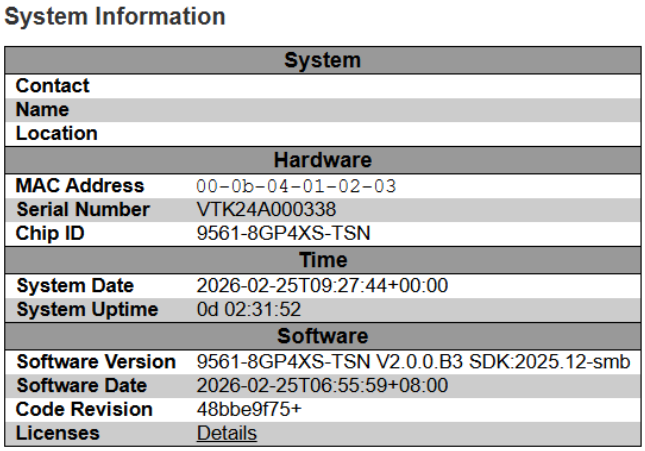

Parameter	Description
Contact	The textual identification of the contact person for this managed node, together with information on how to contact this person as configured in Configuration→System→Information.
Name	An administratively assigned name for this managed node as configured in Configuration→System→Information.
Location	The physical location of this node as configured in Configuration→System→Information.
MAC Address	The MAC address of this switch.
Model Name	The model name of this switch.
System Date	The current (GMT) system time and date. The system time is obtained through the Timing server running on the switch, if any.
System Uptime	The period of time since device has been (re)started.
Software Version	The firmware version currently running on this switch.
Software Date	The build date of the software

Code Revision	The version control identifier of the switch software.
Licenses	Contains version number of the open source

5.1.2 CPU Load

The load is measured as averaged over the last 100ms, 1 sec and 10 seconds intervals. The last 120 samples are graphed, and the last numbers are displayed as text as well.
In order to display the SVG graph, your browser must support the SVG format. Consult the SVG Wiki for more information on browser support. Specifically, at the time of writing, Microsoft Internet Explorer will need to have a plugin installed to support SVG.

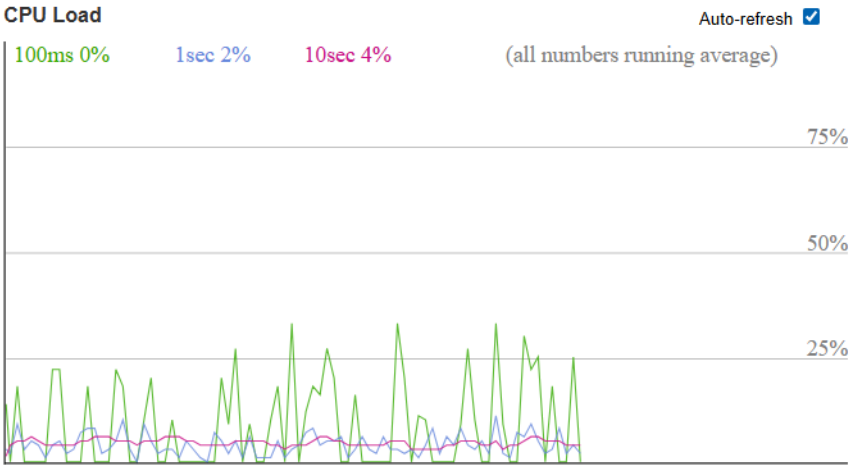

5.1.3 IP Status

This page displays the status of the IP protocol layer. The status is defined by the IP interfaces, the IPv6 routes and the neighbor cache (ARP cache) status.
These settings can be configured in Configuration→System→IP.

IP Interfaces

Interface	Type	Address	Status
VLAN 1	LINK	00-0b-04-14-d8-bf	<UP BROADCAST MULTICAST>
VLAN 1	IPv4	192.168.203.160/24	
VLAN 1	IPv6	fe80::e926:afbb:643c:9fc3/64	

Parameter	Description
Interface	The name of the interface.
Type	The address type of the entry. This may be **LINK**, **IPv4** or **IPv6**.
Address	The current address of the interface (of the given type).
Status	The status flags of the interface (and/or address).

IP Routes

IPv4

Network	Gateway	Status
0.0.0.0/0	192.168.203.1	<UP>
192.168.203.0/24	VLAN 1	<UP>

IPv6

Network	Gateway	Status
fe80::/64	VLAN 1	<UP>

Parameter	Description
Network	The destination IPv4/IPv6 network or host address of this route.
Gateway	The gateway address of this route.
Status	The status flags of the route.

Neighbor cache

IPv4

IP Address	Link Address	State
192.168.203.1	VLAN 1:94-ff-3c-16-6e-91	<REACHABLE HARDWARE>
192.168.203.136	VLAN 1:00-0b-04-14-1d-c7	<STALE HARDWARE>
192.168.203.206	VLAN 1:a0-ad-9f-97-0a-0e	<STALE HARDWARE>
192.168.203.227	VLAN 1:10-c3-7b-97-25-66	<STALE HARDWARE>

IPv6

IP Address	Link Address	State

Parameter	Description
IP Address	The IPv4/IPv6 address of the entry.
Link Address	The Link (MAC) address for which a binding to the IP address given exists.
State	The state of this entry.

5.1.4 IPv4/IPv6 Routing Information Base

Each page shows up to 999 table entries, selected through the "entries per page" input field. When first visited, the web page will show the beginning entries of this table.
The *Start from Network* input field allows the user to change the starting point in this table. Clicking the *Refresh* button will update the displayed table starting from that or the closest next entry match.
In addition, these input fields will upon a *Refresh* button click - assume the value of the first displayed entry, allowing for continuous refresh with the same start input field.

IPv4

Routing Information Base

1 - 2 of 2 entries Auto-refresh

Start from Network 0.0.0.0 / 0 Protocol Static NextHop 192.168.203.1 with 20 entries per page.

Codes: C - connected, S - static, O - OSPF, R - RIP, * - selected route, D - DHCP installed route

Protocol	Network/Prefix	NextHop	Distance	Metric	Interface	Uptime (hh:mm:ss)	State
S *	0.0.0.0/0	192.168.203.1	1	0	VLAN 1	4d 12:44:00	Active
C *	192.168.203.0/24	-	-	-	VLAN 1	4d 12:44:00	Active

IPv6

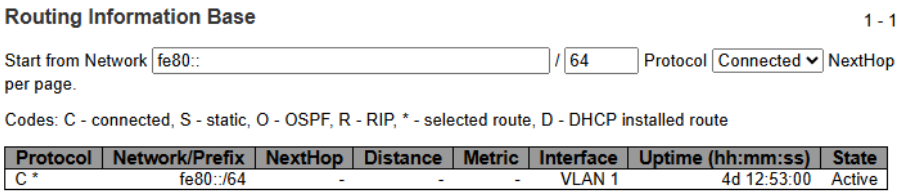

Parameter	Description
Protocol	The protocol that installed this route. **DHCP**: The route is created by DHCP. **Connected**: The destination network is connected directly. **Static**: The route is created by user. **OSPF**: The route is created by OSPF.
Network/Prefix	Network and prefix (example **10.0.0.0/16** or **2000::0/64**) of the given route entry.
NextHop	Next-hop IP address. All-zeroes indicates the link is directly connected, 'blackhole' that it is a blackhole (discarding) route.
Distance	Distance of the route in hops.
Metric	Metric of the route.
Interface	Next-hop interface.
Uptime	Time (hh:mm:ss) since this route was created.
State	Displays the status of the destination.

5.1.5 Log

Each page shows up to 999 table entries, selected through the *entries per page* input field. When first visited, the web page will show the beginning entries of this table.
The *Level* input field is used to filter the display system log entries.
The *Clear Level* input field is used to specify which system log entries will be cleared.
To clear specific system log entries, select the clear level first then click the *Clear* button.
The *Start from ID input* field allows the user to change the starting point in this table. Clicking the *Refresh* button will update the displayed table starting from that or the closest next entry match.
In addition, these input fields will assume the value of the first displayed entry when clicking the *Refresh* button, allowing for continuous refresh with the same start input field.
Logging can be configured at Configuration→System→Log.

VOLKTEK

System Log Information

Auto-refresh ☐ Refresh Clear |<< << >> >>|

Level	All
Clear Level	All

The total number of entries is 108 for the given level.

Start from ID 1 with 20 entries per page.

ID	Level	Time	Message
1	Informational	2025-12-13T15:53:26+00:00	SYS-BOOTING: Switch just made a cold boot (ver 9561-8GP4XS-TSN 1.0.1.B0 SDK:2024.09-smb).
2	Notice	2025-12-13T15:53:26+00:00	LINK-UPDOWN: IP Interface VLAN 1 changed state to down.
3	Notice	2025-12-13T15:53:26+00:00	LINK-UPDOWN: IP Interface VLAN 1 changed state to down.
4	Notice	2025-12-13T15:53:28+00:00	POE-CONTROL: controller found
5	Notice	2025-12-13T15:53:39+00:00	LINK-UPDOWN: Interface GigabitEthernet 1/8, changed state to up.
6	Notice	2025-12-13T15:53:43+00:00	LINK-UPDOWN: IP Interface VLAN 1 changed state to up.
7	Informational	2025-12-15T01:56:07+00:00	User admin logged in as local user (priv level 15)
8	Informational	2025-12-15T08:49:39+00:00	User admin logged in as local user (priv level 15)
9	Informational	2025-12-15T08:49:40+00:00	User admin executed command configure terminal
10	Informational	2025-12-15T08:49:59+00:00	User admin executed command exit
11	Informational	2025-12-15T08:50:02+00:00	User admin executed command show alarm status
12	Informational	2025-12-15T08:50:25+00:00	User admin executed command show alarm status | include sasi
13	Informational	2025-12-15T08:51:18+00:00	User admin executed command show alarm sources | include sasi.status
14	Informational	2025-12-15T08:57:05+00:00	User admin executed command show alarm sources
15	Informational	2025-12-15T09:01:59+00:00	User admin executed command show alarm status
16	Informational	2025-12-15T09:02:25+00:00	User admin executed command show alarm sources | include sasi
17	Informational	2025-12-15T09:03:41+00:00	User admin executed command configure terminal
18	Informational	2025-12-15T09:04:45+00:00	User admin executed command alarm alarm.sasi.NO_RPS sasi.status@NO_RPS==false
19	Informational	2025-12-15T09:04:47+00:00	User admin executed command exit
20	Informational	2025-12-15T09:04:51+00:00	User admin executed command show alarm status

Parameter	Description
ID	The identification integer of the system log entry. Clicking it will lead you to the *Detailed System Log* view.
Level	The level of the system log entry. **`Info`**: The system log entry is belonged information level. **`Warning`**: The system log entry is belonged warning level. **`Error`**: The system log entry is belonged error level.
Time	The system time at which the logged event occurred
Message	The detailed message of the system log entry.

5.1.6 Detailed Log

The Detailed System Log view shows all information for individual log messages, which can be selected by their ID. The parameters are identical to Monitor→System→Log.

Detailed System Log Information

ID	1

Message

Level	Informational
Time	2026-02-24T14:21:18+00:00
Message	SYS-BOOTING: Switch just made a cool boot (ver 9561-8GP4XS-TSN V2.0.0.B2 SDK:2025.12-smb).

5.2 Green Ethernet (Port Power Savings)

This page displays the EEE, ActiPhy and Perfect Reach Port Power Saving Status, which can be configured under Configuration→Green Ethernet→Port Power Savings.

Port Power Savings Status

Port	Link	EEE Cap	EEE Ena	LP EEE Cap	EEE In power save	ActiPhy Savings	PerfectReach Savings
1	●	✕	✕	✕	✕	✕	✕
2	●	✕	✕	✕	✕	✕	✕
3	●	✕	✕	✕	✕	✕	✕
4	●	✕	✕	✕	✕	✕	✕
5	●	✕	✕	✕	✕	✕	✕
6	●	✕	✕	✕	✕	✕	✕
7	●	✕	✕	✕	✕	✕	✕
8	●	✕	✕	✕	✕	✕	✕
9	●	✕	✕	✕	✕	✕	✕
10	●	✕	✕	✕	✕	✕	✕
11	●	✕	✕	✕	✕	✕	✕
12	●	✕	✕	✕	✕	✕	✕

Parameter	Description
Local Port	This is the logical port number for this row.
Link	Shows if the link is up for the port (green = link up, red = link down).
EEE cap	Shows if the port is EEE capable.
EEE Ena	Shows if EEE is enabled for the port (reflects the settings at the Port Power Savings configuration page).
LP EEE cap	Shows if the link partner is EEE capable.
EEE In power save	Shows if the system is currently saving power due to EEE. When EEE is enabled, the system will powered down if no frame has been received or transmitted in 5 uSec.
Actiphy Savings	Shows if the system is currently saving power due to ActiPhy.
PerfectReach Savings	Shows if the system is currently saving power due to PerfectReach.

5.3 Thermal Protection

The Thermal Protection Status page shows the current thermal protection status for switch ports.
The Thermal Protection Port Status table lists each port, the current reported temperature, and whether the port link is operating normally or has been shut down by thermal protection.
Figure: Thermal Protection Status page showing port temperature and thermal port status.

Thermal Protection Status

Thermal Protection Port Status

Port	Temperature		Port status
1	56	°C	Port link operating normally
2	56	°C	Port link operating normally
3	56	°C	Port link operating normally
4	56	°C	Port link operating normally
5	56	°C	Port link operating normally
6	56	°C	Port link operating normally
7	56	°C	Port link operating normally
8	56	°C	Port link operating normally
9	56	°C	Port link operating normally
10	56	°C	Port link operating normally
11	56	°C	Port link operating normally
12	56	°C	Port link operating normally

Parameter	Description
Port	Shows the logical port number in the Thermal Protection Port Status table.
Temperature	Shows the current port temperature reported by the switch, in degrees Celsius.
Port status	Shows the thermal protection state of the port. Port link operating normally indicates that the port is not shut down by thermal protection.

5.4 Ports

This section shows current status information and statistics for all individual RJ45 and SFP ports. The port settings can be changed in Configuration→Ports.

5.4.1 State

The Port Statistics Overview page shows the current link state and traffic counters for each switch port.
Use this page to confirm whether a port link is Up or Down, view the negotiated speed, and check receive/transmit rates, packet counters, byte counters, errors, and drops.
Figure: Port Statistics Overview page showing current port link state and traffic counters.

Port Statistics Overview

Port	Link	Speed	Rate (Kbps)		Packets		Bytes		Errors		Drops	
			Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx
1	Down	-	-	-	0	0	0	0	0	0	0	0
2	Up	1G	0	0	293821225	882002299	433157980878	1295179542600	0	0	0	0
3	Up	1G	0	0	293821611	882001927	433158519363	1295179005165	0	0	0	0
4	Down	-	-	-	0	0	0	0	0	0	0	0
5	Down	-	-	-	0	0	0	0	0	0	0	0
6	Down	-	-	-	0	0	0	0	0	0	0	0
7	Down	-	-	-	0	0	0	0	0	0	0	0
8	Down	-	-	-	0	0	0	0	0	0	0	0
9	Down	-	-	-	0	0	0	0	0	0	0	0
10	Up	1G	0	0	588281384	587617073	862029849105	866344981754	0	0	0	1
11	Down	-	-	-	0	0	0	0	0	0	0	0
12	Down	-	-	-	0	0	0	0	0	0	0	0

Parameter	Description
Port	Shows the logical port number in the Port Statistics Overview table.
Link	Shows the current link state of the port. Up indicates an active link; Down indicates that no active link is detected.
Speed	Shows the current negotiated link speed for an active port. A dash indicates that the port link is down.
Rate (Kbps)	Shows the current receive and transmit traffic rate in Kbps.
Packets	Shows the accumulated receive and transmit packet counters.
Bytes	Shows the accumulated receive and transmit byte counters.
Errors	Shows the accumulated receive and transmit error counters.
Drops	Shows the accumulated receive and transmit drop counters.

5.4.2 Traffic Overview

This table provides an overview of general traffic statistics for all switch ports including errors, dropped and filtered packets.

Port Statistics Overview

Port	Packets		Bytes		Errors		Drops		Filtered
	Received	Transmitted	Received	Transmitted	Received	Transmitted	Received	Transmitted	Received
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	4096650	64619	390829444	39279464	0	0	0	0	1497619
9	0	0	0	0	0	0	0	0	0
10	0	0	0	0	0	0	0	0	0
11	0	0	0	0	0	0	0	0	0
12	0	0	0	0	0	0	0	0	0

Parameter	Description
Port	The logical port number associated with this row.
Packets	The number of received and transmitted packets per port.
Bytes	The number of received and transmitted bytes per port.
Errors	The number of frames received in error and the number of incomplete transmissions per port.
Drops	The number of frames discarded due to ingress or egress congestion.
Filtered	The number of received frames filtered by the forwarding process.

5.4.3 QoS Statistics

This table provides statistics for the different queues for all switch ports. QoS queue settings can be configured at Configuration→QoS→Queue Policing.

Queuing Counters

Port	Q0		Q1		Q2		Q3		Q4		Q5		Q6		Q7	
	Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx
1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
4	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
5	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
6	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
7	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
8	4098265	34068	0	0	0	0	0	0	0	0	0	0	0	0	0	30733
9	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
10	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
11	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
12	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Parameter	Description
Port	The logical port number associated with this row.

Qn	There are 8 QoS queues per port. Q0 is the lowest priority queue.
Rx/Tx	The number of received and transmitted packets per queue.

5.4.4 QCL Status

This table shows the QoS Control List status by different QCL users. Each row describes the QCE that is defined. It is a conflict if a specific QCE is not applied to the hardware due to hardware limitations. The maximum number of QCEs is **256** on each switch.
The QCL can be configured at Configuration→QoS→QoS Control List.

QoS Control List Status

User	QCE	Port	Frame Type	Action							Conflict
				CoS	DPL	DSCP	PCP	DEI	Policy	Ingress Map	
No entries											

Parameter	Description
User	Indicates the QCL user.
QCE	Indicates the QCE id.
Port	Indicates the list of ports configured with the QCE.
Frame Type	Indicates the type of frame. Possible values are: **Any**: Match any frame type. **Ethernet**: Match EtherType frames. **LLC**: Match (LLC) frames. **SNAP**: Match (SNAP) frames. **IPv4**: Match IPv4 frames. **IPv6**: Match IPv6 frames.
Action	Indicates the classification action taken on ingress frame if parameters configured are matched with the frame's content. Possible actions are: **CoS**: Classify Class of Service **DPL**: Classify Drop Precedence Level. **DSCP**: Classify DSCP value. **PCP**: Classify PCP value. **DEI**: Classify DEI value. **Policy**: Classify ACL Policy number. **Ingress Map**: Classify Ingress Map ID.
Conflict	Displays Conflict status of QCL entries. As H/W resources are shared by multiple applications. It may happen that resources required to add a QCE may not be available, in that case it shows conflict status as 'Yes', otherwise it is always 'No'. Please note that conflict can be resolved by releasing the H/W resources required to add QCL entry on pressing 'Resolve Conflict' button
Combined	Select the QCL status from this drop-down list.
Resolve Conflict	Click to release the resources required to add QCL entry, in case the conflict status for any QCL entry is yes.

5.4.5 Detailed Statistics

The detailed statistics section shows RMON statistics for each port individually. Use the port select box to select which switch port details to display.

The displayed counters are the totals for receive and transmit, the size counters for receive and transmit, and the error counters for receive and transmit.

Detailed Port Statistics Port 8

Port 8 | Auto-refresh | Refresh

Receive Total		Transmit Total	
Rx Packets	4643449	Tx Packets	113438
Rx Octets	465871468	Tx Octets	63611655
Rx Unicast	45930	Tx Unicast	63740
Rx Multicast	4267640	Tx Multicast	49695
Rx Broadcast	329879	Tx Broadcast	3
Rx Pause	0	Tx Pause	0
Receive Size Counters		**Transmit Size Counters**	
Rx 64 Bytes	2633001	Tx 64 Bytes	13418
Rx 65-127 Bytes	1589341	Tx 65-127 Bytes	8292
Rx 128-255 Bytes	223616	Tx 128-255 Bytes	48910
Rx 256-511 Bytes	107806	Tx 256-511 Bytes	4241
Rx 512-1023 Bytes	86136	Tx 512-1023 Bytes	5293
Rx 1024-1518 Bytes	3272	Tx 1024-1518 Bytes	33284
Rx 1519- Bytes	277	Tx 1519- Bytes	0
Receive Queue Counters		**Transmit Queue Counters**	
Rx Q0	4643449	Tx Q0	64919
Rx Q1	0	Tx Q1	0
Rx Q2	0	Tx Q2	0
Rx Q3	0	Tx Q3	0
Rx Q4	0	Tx Q4	0
Rx Q5	0	Tx Q5	0
Rx Q6	0	Tx Q6	0
Rx Q7	0	Tx Q7	48519
Receive Error Counters		**Transmit Error Counters**	
Rx Drops	0	Tx Drops	0
Rx CRC/Alignment	0	Tx Late/Exc. Coll.	0
Rx Undersize	0		
Rx Oversize	0		
Rx Fragments	0		
Rx Jabber	0		
Rx Filtered	1616673		
Receive MM Counters		**Transmit MM Counters**	
Rx MM Fragments	0	Tx MM Fragments	0
Rx MM Assembly Ok	0	Tx MM Hold	0
Rx MM Assembly Errors	0		
Rx MM SMD Errors	0		

Receive Total and Transmit Total

Parameter	Description
Rx and Tx Packets	The number of received and transmitted (good and bad) packets.
Rx and Tx Octets	The number of received and transmitted (good and bad) bytes. Includes FCS, but excludes framing bits.
Rx and Tx Unicast	The number of received and transmitted (good and bad) unicast packets.
Rx and Tx Multicast	The number of received and transmitted (good and bad) multicast packets.
Rx and Tx Broadcast	The number of received and transmitted (good and bad) broadcast packets.
Rx and Tx Pause	A count of the MAC Control frames received or transmitted on this port that have an opcode indicating a PAUSE operation.

Receive and Transmit Size Counters

The number of received and transmitted (good and bad) packets split into categories based on their respective frame sizes.

Receive and Transmit Queue Counters

The number of received and transmitted packets per input and output queue.

Receive and Transmit Error Counters

Parameter	Description
Rx Drops	The number of frames dropped due to lack of receive buffers or egress congestion.
Rx CRC/Alignment	The number of frames received with CRC or alignment errors.
Rx Undersize	The number of short[1] frames received with valid CRC.
Rx Oversize	The number of long[2] frames (longer than the configured maximum frame length for this port) received with valid CRC.
Rx Fragments	The number of short[1] frames received with invalid CRC.
Rx Jabber	The number of long[2] frames received with invalid CRC.
Rx Filtered	The number of received frames filtered by the forwarding process.
Tx Drops	The number of frames dropped due to output buffer congestion.
Tx Late/Exc. Coll.	The number of frames dropped due to excessive or late collisions.

[1] Short frames are frames that are smaller than 64 bytes.
[2] Long frames are frames that are longer than the configured maximum frame length for this port.

Receive and Transmit MM Counters

Parameter	Description
Rx MM Fragments	A count of received MAC frame fragments.
Rx MM Assembly Ok	A count of MAC frames that were successfully reassembled and delivered to MAC.
Rx MM Assembly Errors	A count of MAC frames with reassembly errors. The counter is incremented when the ASSEMBLY_ERROR state of the Receive Processing State Diagram is entered.
Rx MM SMD Errors	A count of received MAC frames / MAC frame fragments rejected due to unknown SMD value or arriving with an SMD-C when no frame is in progress. The counter is incremented each time the BAD_FRAG state of the Receive Processing State Diagram is entered
Tx MM Fragments	A count of transmitted MAC frame fragments.
Tx MM Hold	A count of times MM_CTL.request (HOLD) primitive assertion caused preemption of a preemptable MAC frame.

5.4.6 Name Map

Many Web pages use a port number to express an interface, whereas CLI uses interface names. The table on this page provides a means to convert from one to the other.

Interface Name to Port Number Map

Interface Name	Port Number
Gi 1/1	1
Gi 1/2	2
Gi 1/3	3
Gi 1/4	4
Gi 1/5	5
Gi 1/6	6
Gi 1/7	7
Gi 1/8	8
10G 1/1	9
10G 1/2	10
10G 1/3	11
10G 1/4	12

5.5 CFM

Monitor the status of the CFM MEPs.
CFM MEPs can be configured at Configuration→CFM.

CFM MEP Status

Domain	Service	MEPID	Port	State		SMAC	Defects		CCM Rx			CCM Tx
				Active	Fng		Highest	Defects	Valid	Invalid	Errors	
No entry exists												

Parameter	Description
Domain	Name of Domain under which this MEP resides.
Service	Name of Service under which this MEP resides.
MEPID	The identification of this MEP.
Port	Port on which this MEP resides.
State	*Active*: Operational state of the MEP. : OFF. This indicates that the MEP Admin State is disabled. : DOWN. The MEP Admin State is enabled, but an error state exists. : UP. The MEP Admin State is enabled, and no errors and defects exists. *Fng*: Holds the current state of the Fault Notification Generator State Machine. Values will be one of the following: **State** — **Description** reset — No defect has been present since reset timer expired or the State Machine was last reset. defect — A defect is present, but not for a long enough time to be reported. reportDefect — A transient state during which the defect is reported. defectReported — A defect is present, and some defect has been reported. defectClearing — No defect is present, but the ResetTime timer has not yet expired.
SMAC	This MEP's MAC address.
Defects	*Highest*: Highest priority defect that has been present since the MEP's fault notification generator state machine was last in the reset state. *Defects*: A MEP can detect and report a number of defects, and multiple defects can be present at the same time. This is indicated the following letter code. **Code** — **Defect** — **Description** - — Defect not present — Defect not present R — someRDIdefect — RDI received from at least one remote MEP M — someMACstatusDefect — Received Port Status TLV != psUp or Interface Status TLV != isUp C — someRMEPCCMdefect — Valid CCM is not received within 3.5 times CCM interval from at least one remote MEP E — errorCCMdefect — Received CCM from an unknown remote MEP-ID or CCM interval mismatch X — xconCCMdefect — Received CCM with an MD/MEG level smaller than configured or wrong MAID/MEGID (cross-connect)
CCM Rx	*Valid*: Total number of CCMs that hit this MEP and passed the validation test. *Invalid*: Total number of CCMs that hit this MEP and didn't pass the validation test. *Errors*: Total number of out-of-sequence errors seen from RMEPs.
CCM Tx	Total number of CCM PDUs transmitted by this MEP.

5.6 APS

This shows the current status of the APS instances. APS can be configured at Configuration→APS.

APS Status

APS #	State			Defect state		TxAps			RxAps			Dfop				SMAC	TxCnt	RxCnt	
	Operational	Warning	Protection	Working	Protecting	Request	ReSignal	BrSignal	Request	ReSignal	BrSignal	CM	PM	NR	TO			Valid	Invalid
No entry exists																			

Parameter	Description
APS #	The ID of the APS. Click on link to get to APS instance page, you can reset counters and issue commands.
State, Operational	The operational state of the APS instance. There are many ways to not have the instance active. Each of them has its own value. Only when the state is Active, the APS instance will be active and up and running. If the Operational state is not "Active", the remaining fields are invalid. The possible values of this field are shown below: **Administratively disabled** **Active** **Internal Error** **Working MEP not Found** **Protecting MEP not Found** **Working MEP is not administrative active** **Protecting MEP is not administrative active** **Working MEP is not a Down MEP** **Protecting MEP is not a Down MEP** **Working and Protecting MEP use the same interface** **Another instance use the same Working port**
State, Warning	If the operational state is Active, the APS instance is indeed active, but it may be that it doesn't run as the administrator thinks, because of configuration errors, which are reflected in the warnings below. The Warning information is indicated by ●: no warning, ●: warning. Use the tooltip to get the detailed warning information.
State, Protection	The possible protection group states. The letters refer to the state as described in G.8031 Annex **No request Working:** A. **No request Protecting:** B. **Lockout:** C. **Forced Switch:** D. **Signal fail Working:** E. **Signal fail Protecting:** F. **Manual switch to Protecting:** G. **Manual switch to Working:** H. **Wait to restore:** I. **Do not revert:** J. **Exercise Working:** K. **Exercise Protecting:** L. **Reverse request Working:** M. **Reverse request Protecting:** N. **Signal degrade Working:** P. **Signal degrade Protecting:** Q.
Defect state, Working/Protection	The possible values of this field are shown below: **ok:** The port defect state is OK **sd:** The port defect state is Signal Degrade **sf:** The port defect state is Signal Fail
TxAps/RxAps, Request	The possible transmitted or received APS request according to G.8031, Table 11-1. **nr:** No Request. **dnr:** Do Not Revert. **rr:** Reverse Request. **exer:** Exercise. **wtr:** Wait-To-Restore. **ms:** Manual Switch. **sd:** Signal Degrade. **sfW:** Signal Fail for Working. **fs:** Forced Switch. **sfP:** Signal Fail for Protect.

	lo: Lockout.
TxAps, ReSignal	Transmitted requested signal according to G.8031 figure 11-2
TxAps, BrSignal	Transmitted bridged signal according to G.8031 figure 11-2
RxAps, ReSignal	Received requested signal according to G.8031 figure 11-2
RxAps, BrSignal	Received bridged signal according to G.8031 figure 11-2
Dfop	Dfop is "Failure of Protocol defect" and the presence of a defect is indicated by 🟢: no defect, 🔴: defect. **CM:** Configuration Mismatch (received APS PDU on working interface within last 17.5 seconds). **PM:** Provisioning Mismatch (far and near ends are not using the same mode; bidir only) **NR:** No Response (far end hasn't agreed on 'Requested Signal' within 50 ms; bidir only) **TO:** Time Out (near end hasn't received a valid APS PDU within last 17.5 seconds; bidir only)
SMAC	Source MAC address of last received APS PDU or all-zeros if no PDU has been received.
TxCnt	Number of APS PDU frames transmitted.
RxCnt, Valid	Number of valid APS PDU frames received on the protect port.
RxCnt, Invalid	Number of invalid APS PDU frames received on the protect port.

5.7 ERPS

This shows the current status of the ERPS instances. ERPS can be configured at Configuration→ERPS.ERPS can be configured at Configuration→ERPS.

ERPS Status

ERPS #	Oper	Warning	State	TxRapsActive	cFOPTo	Tx Info							
						Update Time Secs	Request	Version	Rb	Dnf	Bpr	Node Id	SMAC
No entry exists													

Parameter	Description
ERPS #	The ID of the ERPS. Click on link to get to ERPS detailed instance page, you can reset counters and issue commands.
Oper	The operational state of ERPS instance. 🟢: Active. 🔴: Disabled or Internal error
Warning	Operational warnings of ERPS instance. ⚫: No warnings. 🟡: There are warnings, use tooltip to see
State	Specifies protection/node state of ERPS.
TxRapsActive	Specifies whether we are currently supposed to be transmitting R-APS PDUs on our ring ports.
cFOPTo	Failure of Protocol - R-APS Rx Time Out.
Update Time Secs	Time in seconds since boot that this structure was last updated.
Request	Request/state according to G.8032, table 10-3.
Version	Version of received/used R-APS Protocol. 0 means v1, 1 means v2, etc.
Rb	RB (RPL blocked) bit of R-APS info. See Figure 10-3 of G.8032.

Dnf	DNF (Do Not Flush) bit of R-APS info. See Figure 10-3 of G.8032.
Bpr	BPR (Blocked Port Reference) of R-APS info. See Figure 10-3 of G.8032.
Node Id	Node ID of this request.
SMAC	The Source MAC address used in the request/state.

5.8 Media Redundancy Protocol

This shows the current status of the MRP instances.
Media Redundancy Protocol can be configured at Configuration→Media Redundancy.

Media Redundancy Protocol Status Overview

Inst #	Opererational			Ring State	Interconnection State
	State	Warnings	Role		
1			Client	Disabled	Disabled

Parameter	Description
Inst #	The ID of the MRP instance. Click on link to navigate to the detailed MRP status page, where you can see more status and statistics.
Operational State	Indicates the operational state of the manager for the entire ring: **Green**: Active (forwarding/blocking as expected). **Red**: Stopped, recovery, or internal failure.
Operational Warnings	Indicates if any operational warnings (operwarnings) are active: **Yellow**: Warnings exist (e.g., VLAN tagging mismatches, signal fail fallbacks, multiple managers). **Gray**: No warnings. The indicator is yellow if the MRP instance is enabled and contains operational warnings. In that case, hover the mouse over the yellow image and wait for a list of warnings to appear.
Operational Role	Shows the current role of the MRP instance. This is particularly useful if configured as an MRA, where it can assume both the role of an MRM and an MRC.
Ring State	Reports the current state of the ring: **Closed**: The ring is healthy and complete. The MRM blocks one port to prevent loops. **Open**: The ring is broken (e.g., link failure). The MRM unblocks ports to restore connectivity.
Interconnection State	Reports the state of the interconnection ring (if configured): **Closed**: The interconnection ring is active and healthy. **Open**: The interconnection ring is broken or in an open state.

The following information provides detailed status and statistics of a given MRP instance.
This shows the current configuration for this MRP instance. For an explanation of the individual fields, refer to Configuration→Media Redundancy page.

Media Redundancy Protocol Detailed Status

Auto-refresh ☐ Refresh

Configuration

Inst #	Enabled	Ring								Interconnection					
		Role	Name	Port1		Port2		VLAN	Recovery Profile	UUID	Role	Name	Port	SF	VLAN	Recovery Profile
				Port	SF	Port	SF									
1	✗	Auto Manager		none	Link	none	Link	0	500 ms	ffffffffffffffffffffffffffffffff	None		none	Link	0	500 ms

This shows the current status of the MRP instance.

Status

Operational			Flushes	Ring						Interconnection				
State	Warnings	Role		State	Transitions		Round-trip Time (ms)			State	Transitions	Round-trip Time (ms)		
					Open/Closed	MRM/MRC	Min	Last	Max		Open/Closed	Min	Last	Max
		Client	0	Disabled	0	0	-	-	-	Disabled	0	-	-	-

Parameter	Description
Operational State	Indicates the operational state of the manager for the entire ring: **Green**: Active (forwarding/blocking as expected). **Red**: Stopped, recovery, or internal failure.
Operational Warnings	Indicates if any operational warnings (operwarnings) are active: **Yellow**: Warnings exist (e.g., VLAN tagging mismatches, signal fail fallbacks, multiple managers). **Gray**: No warnings. The indicator is yellow if the MRP instance is enabled and contains operational warnings. In that case, hover the mouse over the yellow image and wait for a list of warnings to appear.
Operational Role	Shows the current role of the MRP instance. This is particularly useful if configured as an MRA, where it can assume both the role of an MRM and an MRC.
Flushes	Shows the number of times the Forwarding Database (FDB) has been flushed.
Ring State	Reports the current state of the ring: **Closed**: The ring is healthy and complete. The MRM blocks one port to prevent loops. **Open**: The ring is broken (e.g., link failure). The MRM unblocks ports to restore connectivity.
Ring Transitions Open/Closed	Shows the number of times the ring has gone from being open to being closed and vice versa. It will always be 0 on nodes operating as MRCs.
Ring Transitions MRM/MRC	Shows the number of times that this node has gone from being an MRM to becoming an MRC or vice versa. Only relevant if the configured role is MRA.
Ring Round-trip Time (Min/Last/Max)	Shows the number of milliseconds an MRP_Test PDU transmitted by this node spent in the ring until it came back. The minimum and maximum times are shown along with the time it took for the last MRP_Test PDU to travel the ring. Only valid if the MRP instance is currently operating as an MRM and if not using hardware to transmit MRP_Test PDUs.
Interconnection State	Reports the state of the interconnection ring (if configured): **Closed**: The interconnection ring is active and healthy. **Open**: The interconnection ring is broken or in an open state.
Interconnection Transitions Open/Closed	Shows the number of times the interconnection topology has gone from being open to a closed topology. It only counts on the MIM.
Interconnection Round-trip Time (Min/Last/Max)	Shows the number of milliseconds an MRP_InTest PDU transmitted by this node spent in the interconnection ring until it came back. The minimum and maximum times are shown along with the time it took for the last MRP_InTest PDU to travel the ring. Only valid if the MRP instance is a MIM and if not using hardware to transmit MRP_InTest PDUs.

This shows statistics collected on the MRP instance.

Statistics

Counter	Rx			Tx		
	Port1	Port2	Interconnection	Port1	Port2	Interconnection
Test	0	0	-	0	0	-
TopologyChange	0	0	-	0	0	-
LinkDown	0	0	-	0	0	-
LinkUp	0	0	-	0	0	-
TestMgrNAck	0	0	-	0	0	-
TestPropagate	0	0	-	0	0	-
Option	0	0	-	0	0	-
InTest	0	0	-	0	0	-
InTopologyChange	0	0	-	0	0	-
InLinkDown	0	0	-	0	0	-
InLinkUp	0	0	-	0	0	-
InLinkStatusPoll	0	0	-	0	0	-
Unknown	0	0	-			
Errors	0	0	-			
Unhandled	0	0	-			
Own	0	0	-			
Signal Fails	0	0	-			

Parameter	Description
Rx/Tx Test	Shows - per port - the number of received and transmitted MRP_Test PDUs. On switches with hardware support for transmitting MRP_Test PDUs, the Tx counters will not count, because the actual number of transmitted PDUs cannot be obtained.
Rx/Tx TopologyChange	Shows - per port - the number of received and transmitted MRP_TopologyChange PDUs.
Rx/Tx LinkDown	Shows - per port - the number of received and transmitted MRP_LinkDown PDUs.
Rx/Tx LinkUp	Shows - per port - the number of received and transmitted MRP_LinkUp PDUs.
Rx/Tx TestMgrNAck	Shows - per port - the number of received and transmitted MRP_TestMgrNAck PDUs.
Rx/Tx TestPropagate	Shows - per port - the number of received and transmitted MRP_TestPropagate PDUs.
Rx/Tx Option	Shows - per port - the number of received and transmitted MRP_Option PDUs.
Rx/Tx In Test	Shows - per port - the number of received and transmitted MRP_InTest PDUs. On switches with hardware support for transmitting MRP_InTest PDUs, the Tx counters will not count, because the actual number of transmitted PDUs cannot be obtained.
Rx/Tx In ToppologyChange	Shows - per port - the number of received and transmitted MRP_InTopologyChange PDUs.
Rx/Tx InLinkDown	Shows - per port - the number of received and transmitted MRP_InLinkDown PDUs.
Rx/Tx InLinkUp	Shows - per port - the number of received and transmitted MRP_InLinkUp PDUs.
Rx/Tx InLinkStatusPoll	Shows - per port - the number of received and transmitted MRP_InLinkStatusPoll PDUs.
Unknown	Shows - per port - the number of received MRP PDUs that wasn't recognized as a known MRP_Xxx PDU type.
Errors	Shows - per port - the number of received erroneous MRP PDUs. An erroneous MRP PDU could e.g. be if the frame is too short to carry all the required fields of a given MRP PDU type.
Unhandled	Shows - per port - the number of MRP PDUs received by, but not processed by, the node for one or the other reason.
Own	Shows - per port - the number of MRP PDUs that were transmitted by the node itself.
Signal Fails	Shows - per port - the number of signals fails that have occurred.

5.9 Link OAM

OAM is an acronym for **O**peration **A**dministration and **M**aintenance.

5.9.1 Statistics

This shows detailed OAM traffic statistics for a specific switch port. Use the port select box to select which switch port details to display.

The displayed counters represent the total number of OAM frames received and transmitted for the selected port. Discontinuities of these counters can occur at re-initialization of the management system.
Link OAM can be configured at Configuration→Link OAM.

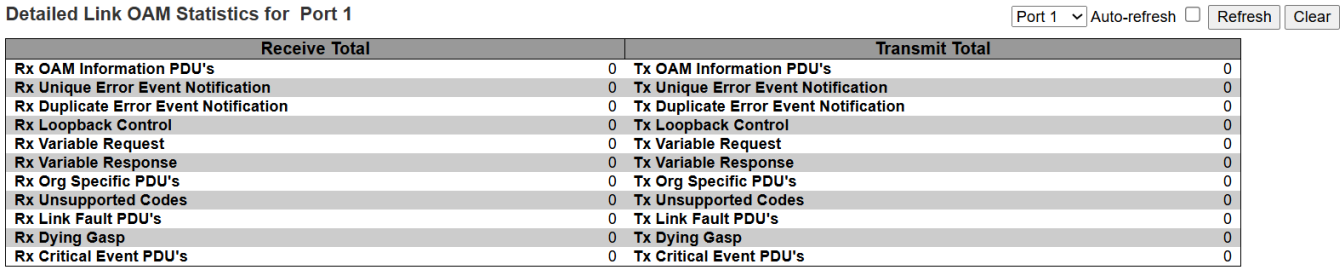

Parameter	Description
Rx and Tx OAM Information PDU's	The number of received and transmitted OAM Information PDU's. Discontinuities of this counter can occur at re-initialization of the management system.
Rx and Tx Unique Error Event Notification	A count of the number of unique Event OAMPDUs received and transmitted on this interface. Event Notifications may be sent in duplicate to increase the probability of successfully being received, given the possibility that a frame may be lost in transit. Duplicate Event Notification transmissions are counted by Duplicate Event Notification counters for Tx and Rx respectively. A unique Event Notification OAMPDU is indicated as an Event Notification OAMPDU with a Sequence Number field that is distinct from the previously transmitted Event Notification OAMPDU Sequence Number.
Rx and Tx Duplicate Error Event Notification	A count of the number of duplicate Event OAMPDUs received and transmitted on this interface. Event Notification OAMPDUs may be sent in duplicate to increase the probability of successfully being received, given the possibility that a frame may be lost in transit. A duplicate Event Notification OAMPDU is indicated as an Event Notification OAMPDU with a Sequence Number field that is identical to the previously transmitted Event Notification OAMPDU Sequence Number.
Rx and Tx Loopback Control	A count of the number of Loopback Control OAMPDUs received and transmitted on this interface.
Rx and Tx Variable Request	A count of the number of Variable Request OAMPDUs received and transmitted on this interface.
Rx and Tx Variable Response	A count of the number of Variable Response OAMPDUs received and transmitted on this interface.
Rx and Tx Org Specific PDU's	A count of the number of Organization Specific OAMPDUs transmitted on this interface.
Rx and Tx Unsupported Codes	A count of the number of OAMPDUs transmitted on this interface with an unsupported op-code.
Rx and Tx Link fault PDU's	A count of the number of Link fault PDU's received and transmitted on this interface.
Rx and Tx Dying Gasp	A count of the number of Dying Gasp events received and transmitted on this interface.
Rx and Tx Critical Event PDU's	A count of the number of Critical event PDU's received and transmitted on this interface.

5.9.2 Port Status

This shows Link OAM configuration operational status.
The displayed fields shows the active configuration status for the selected port.
Link OAM can be configured at Configuration→Link OAM.

Detailed Link OAM Status for

Port 1 Auto-refresh Refresh

PDU Permission	Receive Only
Discovery State	Fault State
Peer MAC Address	00:00:00:00:00:00

Local		Peer	
Mode	Disabled	**Mode**	Disabled
Unidirectional Operation Support	Disabled	**Unidirectional Operation Support**	0
Remote Loopback Support	Disabled	**Remote Loopback Support**	Disabled
Link Monitoring Support	Disabled	**Link Monitoring Support**	Disabled
MIB Retrieval Support	Disabled	**MIB Retrieval Support**	Disabled
MTU Size	1516	**MTU Size**	0
Multiplexer State	Discarding	**Multiplexer State**	Discarding
Parser State	Discarding	**Parser State**	Discarding
Organizational Unique Identification	00-00-00	**Organizational Unique Identification**	00-00-00
PDU Revision	0	**PDU Revision**	0

Parameter	Description
Mode	The Mode in which the Link OAM is operating, Active or Passive.
Unidirectional Operation Support	This feature is not available to be configured by the user. The status of this configuration is retrieved from the PHY.
Remote Loopback Support	If status is enabled, DTE is capable of OAM remote loopback mode.
Link Monitoring Support	If status is enabled, DTE supports interpreting Link Events.
MIB Retrieval Support	If status ie enabled DTE supports sending Variable Response OAMPDUs.
MTU Size	It represents the largest OAMPDU, in octets, supported by the DTE. This value is compared to the remote's Maximum PDU Size and the smaller of the two is used.
Multiplexer State	When in forwarding state, the Device is forwarding non-OAMPDUs to the lower sublayer. In case of discarding, the device discards all the non-OAMPDU's.
Parser State	When in forwarding state, Device is forwarding non-OAMPDUs to higher sublayer. When in loopback, Device is looping back non-OAMPDUs to the lower sublayer. When in discarding state, Device is discarding non-OAMPDUs.
Organizational Unique Identification	24-bit Organizationally Unique Identifier of the vendor.
PDU Revision	It indicates the current revision of the Information TLV. The value of this field shall start at zero and be incremented each time something in the Information TLV changes. Upon reception of an Information TLV from a peer, an OAM client may use this field to decide if it needs to be processed (an Information TLV that is identical to the previous Information TLV doesn't need to be parsed as nothing in it has changed).
PDU Permission	This field is available only for the Local DTE. It displays the current permission rules set for the local DTE. Possible values are: **`Link fault`**, **`Receive only`**, **`Information exchange only`**, or **`ANY`**.
Discovery State	Displays the current state of the discovery process. Possible states are: **`Fault state`**, **`Active state`**, **`Passive state`**, **`SEND_LOCAL_REMOTE_STATE`**, **`SEND_LOCAL_REMOTE_OK_STATE`**, **`SEND_ANY_STATE`**.

5.9.3 Event Status

This page allows the user to inspect the current Link OAM Link Event configurations and change them as well. The left pane displays the Event status for the Local OAM unit while the right pane displays the status for the Peer for the respective port. Link OAM can be configured at Configuration→Link OAM→Event Settings.

VOLKTEK

Detailed Link OAM Link Status for Port 1

Port 1 Auto-refresh

Local Frame Error Status		Remote Frame Error Status	
Sequence Number	0		
Frame Error Event Timestamp	0	Frame Error Event Timestamp	0
Frame error event window	0	Frame error event window	0
Frame error event threshold	0	Frame error event threshold	0
Frame errors	0	Frame errors	0
Total frame errors	0	Total frame errors	0
Total frame error events	0	Total frame error events	0
Local Frame Period Status		**Remote Frame Period Status**	
Frame Period Error Event Timestamp	0	Frame Period Error Event Timestamp	0
Frame Period Error Event Window	0	Frame Period Error Event Window	0
Frame Period Error Event Threshold	0	Frame Period Error Event Threshold	0
Frame Period Errors	0	Frame Period Errors	0
Total frame period errors	0	Total frame period errors	0
Total frame period error events	0	Total frame period error events	0
Local Symbol Period Status		**Remote Symbol Period Status**	
Symbol Period Error Event Timestamp	0	Symbol Period Error Event Timestamp	0
Symbol Period Error Event Window	0	Symbol Period Error Event Window	0
Symbol Period Error Event Threshold	0	Symbol Period Error Event Threshold	0
Symbol Period Errors	0	Symbol Period Errors	0
Total symbol period errors	0	Total symbol period errors	0
Total Symbol period error events	0	Total Symbol period error events	0
Local Event Seconds Summary Status		**Remote Event Seconds Summary Status**	
Error Frame Seconds Summary Event Timestamp	0	Error Frame Seconds Summary Event Timestamp	0
Error Frame Seconds Summary Event window	0	Error Frame Seconds Summary Event window	0
Error Frame Seconds Summary Event Threshold	0	Error Frame Seconds Summary Event Threshold	0
Error Frame Seconds Summary Errors	0	Error Frame Seconds Summary Errors	0
Total Error Frame Seconds Summary Errors	0	Total Error Frame Seconds Summary Errors	0
Total Error Frame Seconds Summary Events	0	Total Error Frame Seconds Summary Events	0

Parameter	Description
Port	The switch port number.
Sequence Number	This two-octet field indicates the total number of events occurred at the remote end.
Frame Error Event Timestamp	This two-octet field indicates the time reference when the event was generated, in terms of 100 ms intervals.
Frame error event window	This two-octet field indicates the duration of the period in terms of 100 ms intervals. 1) The default value is one second. 2) The lower bound is one second. 3) The upper bound is one minute.
Frame error event threshold	This four-octet field indicates the number of detected errored frames in the period is required to be equal to or greater than in order for the event to be generated. 1) The default value is one frame error. 2) The lower bound is zero frame errors. 3) The upper bound is unspecified.
Frame errors	This four-octet field indicates the number of detected errored frames in the period.
Total frame errors	This eight-octet field indicates the sum of errored frames that have been detected since the OAM sublayer was reset.
Total frame error events	This four-octet field indicates the number of Errored Frame Event TLVs that have been generated since the OAM sublayer was reset.
Frame Period Error Event Timestamp	This two-octet field indicates the time reference when the event was generated, in terms of 100 ms intervals.
Frame Period Error Event Window	This four-octet field indicates the duration of period in terms of frames.
Frame Period Error Event Threshold	This four-octet field indicates the number of errored frames in the period is required to be equal to or greater than in order for the event to be generated.
Frame Period Errors	This four-octet field indicates the number of frame errors in the period.
Total frame period errors	This eight-octet field indicates the sum of frame errors that have been detected since the OAM sublayer was reset.
Total frame period error events	This four-octet field indicates the number of Errored Frame Period Event TLVs that have been generated since the OAM sublayer was reset.
Symbol Period Error Event Timestamp	This two-octet field indicates the time reference when the event was generated, in terms of 100 ms intervals.
Symbol Period Error Event Window	This eight-octet field indicates the number of symbols in the period.

Symbol Period Error Event Threshold	This eight-octet field indicates the number of errored symbols in the period is required to be equal to or greater than in order for the event to be generated.
Symbol Period Errors	This eight-octet field indicates the number of symbol errors in the period.
Total symbol period errors	This eight-octet field indicates the sum of symbol errors since the OAM sublayer was reset.
Total Symbol period error events	This four-octet field indicates the number of Errored Symbol Period Event TLVs that have been generated since the OAM sublayer was reset.
Error Frame Seconds Summary Event Timestamp	This two-octet field indicates the time reference when the event was generated, in terms of 100 ms intervals, encoded as a 16-bit unsigned integer.
Error Frame Seconds Summary Event window	This two-octet field indicates the duration of the period in terms of 100 ms intervals, encoded as a 16-bit unsigned integer.
Error Frame Seconds Summary Event Threshold	This two-octet field indicates the number of errored frame seconds in the period is required to be equal to or greater than in order for the event to be generated, encoded as a 16-bit unsigned integer.
Error Frame Seconds Summary Errors	This two-octet field indicates the number of errored frame seconds in the period, encoded as a 16-bit unsigned integer.
Total Error Frame Seconds Summary Errors	This four-octet field indicates the sum of errored frame seconds that have been detected since the OAM sublayer was reset.
Total Error Frame Seconds Summary Events	This four-octet field indicates the number of Errored Frame Seconds Summary Event TLVs that have been generated since the OAM sublayer was reset, encoded as a 32bit unsigned integer.

5.10 DHCPv4

DHCP is an acronym for Dynamic Host Configuration Protocol. It is a protocol used for assigning dynamic IP addresses to devices on a network.

5.10.1 Server

DHCP Server is used to allocate network addresses and deliver configuration parameters to dynamically configured hosts called DHCP client.

5.10.1.1 Statistics

This shows the database counters and the number of DHCP messages sent and received by DHCP server. DHCPv4 Server can be configured at Configuration→DHCPv4→Server.

Database Counters

Pool	Excluded IP Address	Declined IP Address
2	2	0

Parameter	Description
Pool	Number of pools.

Excluded IP Address	Number of excluded IP address ranges.
Declined IP Address	Number of declined IP addresses.

Binding Counters

Automatic Binding	Manual Binding	Expired Binding
0	0	0

Parameter	Description
Automatic Binding	Number of bindings with network-type pools.
Manual Binding	Number of bindings that administrator assigns an IP address to a client. That is, the pool is of host type.
Expired Binding	Number of bindings that their lease time expired or they are cleared from Automatic/Manual type bindings.

DHCP Message Received Counters

DISCOVER	REQUEST	DECLINE	RELEASE	INFORM
0	0	0	0	0

Parameter	Description
DISCOVER	Number of DHCP DISCOVER messages received.
REQUEST	Number of DHCP REQUEST messages received.
DECLINE	Number of DHCP DECLINE messages received.
RELEASE	Number of DHCP RELEASE messages received.
INFORM	Number of DHCP INFORM messages received.

DHCP Message Sent Counters

OFFER	ACK	NAK
0	0	0

Parameter	Description
OFFER	Number of DHCP OFFER messages sent.
ACK	Number of DHCP ACK messages sent.
NAK	Umber of DHCP NAK messages sent.

5.10.1.2 Binding

This shows bindings generated for DHCP clients. DHCPv4 Server can be configured at Configuration→DHCPv4→Server.

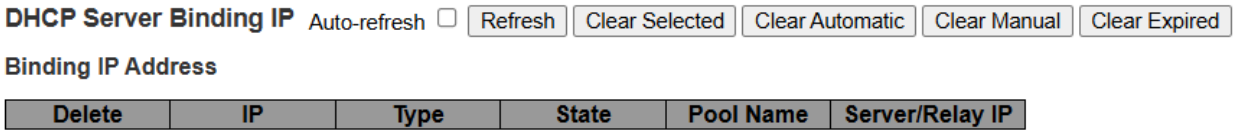

Parameter	Description
IP	IP address allocated to DHCP client.
Type	Type of binding. Possible types are Automatic, Manual, Expired.
State	State of binding. Possible states are Committed, Allocated, Expired.
Pool Name	The pool that generates the binding.
Server/Relay IP	Either IP address of dhcp server or, in case of relayed binding, IP address of relay agent through which binding was negotiated.

5.10.1.3 Declined IP

This shows declined IP address.
DHCPv4 Server can be configured at Configuration→DHCPv4→Server.

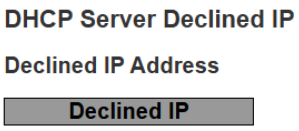

Parameter	Description
Declined IP	List of IP addresses declined.

5.10.2 Snooping Table

This shows the dynamic IP assigned information after DHCP Snooping mode is disabled. All DHCP clients obtained the dynamic IP address from the DHCP server will be listed in this table except for local VLAN interface IP address. Entries in the Dynamic DHCP snooping Table are shown on this page.
DHCPv4 Snooping can be configured at Configuration→DHCPv4→Snooping.

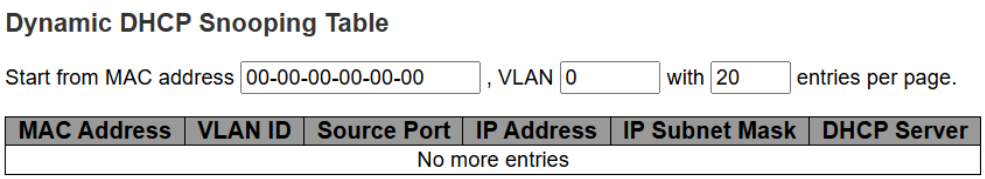

Parameter	Description
MAC Address	User MAC address of the entry.
VLAN ID	VLAN-ID in which the DHCP traffic is permitted.
Source Port	Switch Port Number for which the entries are displayed.
IP Address	User IP address of the entry.

IP Subnet Mask	User IP subnet mask of the entry.
DHCP Server Address	DHCP Server address of the entry.

5.10.3 Relay Statistics

This shows statistics for DHCP relay. DHCP Relay is used to forward and to transfer DHCP messages between the clients and the server when they are not on the same subnet domain.
DHCPv4 relay can be configured at Configuration→DHCPv4→Relay.

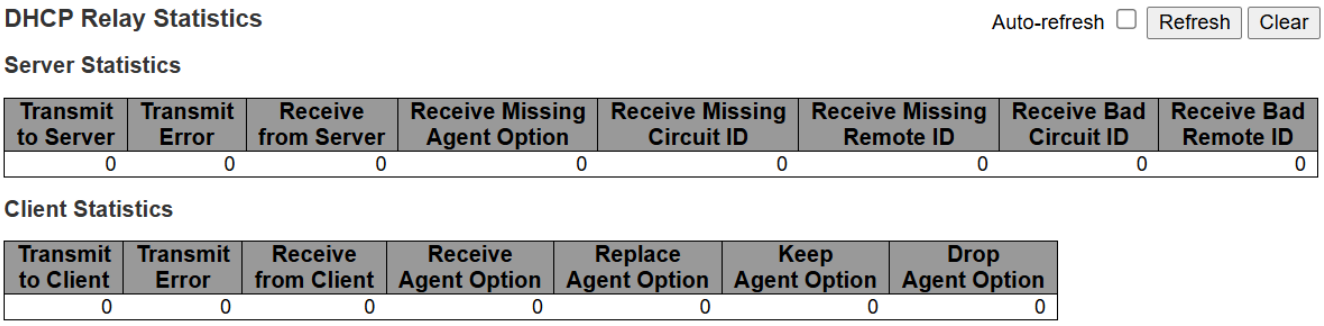

Parameter	Description
Transmit to Server	The number of packets that are relayed from client to server.
Transmit Error	The number of packets that resulted in errors while being sent to clients.
Receive from Server	The number of packets received from server.
Receive Missing Agent Option	The number of packets received without agent information options.
Receive Missing Circuit ID	The number of packets received with the Circuit ID option missing.
Receive Missing Remote ID	The number of packets received with the Remote ID option missing.
Receive Bad Circuit ID	The number of packets whose Circuit ID option did not match known circuit ID.
Receive Bad Remote ID	The number of packets whose Remote ID option did not match known Remote ID.
Transmit to Client	The number of relayed packets from server to client.
Transmit Error	The number of packets that resulted in error while being sent to servers.
Receive from Client	The number of received packets from server.
Receive Agent Option	The number of received packets with relay agent information option.
Replace Agent Option	The number of packets which were replaced with relay agent information option.
Keep Agent Option	The number of packets whose relay agent information was retained.
Drop Agent Option	The number of packets that were dropped which were received with relay agent information.

5.10.4 Detailed Statistics

This shows statistics for DHCP snooping. DHCP Snooping is used to block intruder on the untrusted ports of the switch device when it tries to intervene by injecting a bogus DHCP reply packet to a legitimate conversation between the DCHP client and server. Notice that the normal forward per-port TX statistics isn't increased if the incoming DHCP packet is done by L3 forwarding mechanism. And clear the statistics on specific port may not take effect on global statistics since it gathers the different layer overview.
DHCPv4 Snooping can be configured at Configuration→DHCPv4→Snooping.

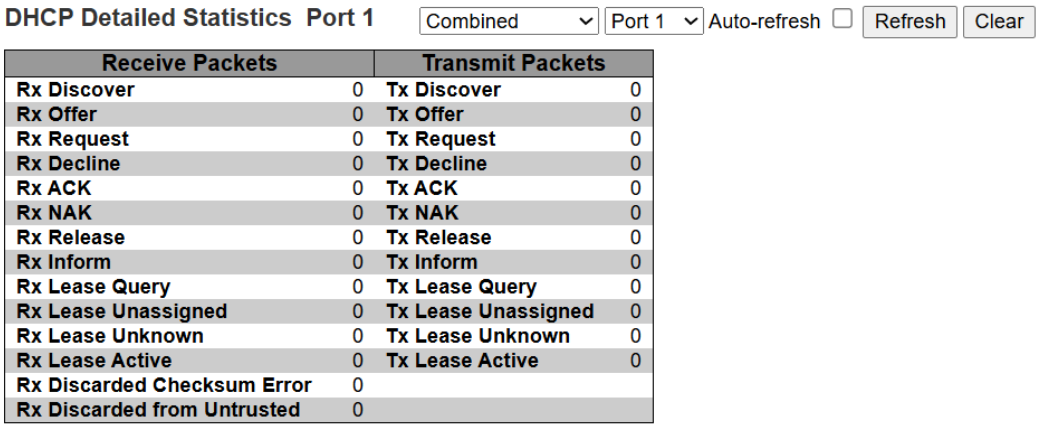

Parameter	Description
Rx and Tx Discover	The number of discover (option 53 with value 1) packets received and transmitted.
Rx and Tx Offer	The number of offer (option 53 with value 2) packets received and transmitted.
Rx and Tx Request	The number of request (option 53 with value 3) packets received and transmitted.
Rx and Tx Decline	The number of decline (option 53 with value 4) packets received and transmitted.
Rx and Tx ACK	The number of ACK (option 53 with value 5) packets received and transmitted.
Rx and Tx NAK	The number of NAK (option 53 with value 6) packets received and transmitted.
Rx and Tx Release	The number of release (option 53 with value 7) packets received and transmitted.
Rx and Tx Inform	The number of inform (option 53 with value 8) packets received and transmitted.
Rx and Tx Lease Query	The number of lease query (option 53 with value 10) packets received and transmitted.
Rx and Tx Lease Unassigned	The number of lease unassigned (option 53 with value 11) packets received and transmitted.
Rx and Tx Lease Unknown	The number of lease unknown (option 53 with value 12) packets received and transmitted.
Rx and Tx Lease Active	The number of lease active (option 53 with value 13) packets received and transmitted.
Rx Discarded checksum error	The number of discard packet that IP/UDP checksum is error.
Rx Discarded from Untrusted	The number of discarded packets coming from untrusted port.

5.11 DHCPv6

5.11.1 Snooping Table

This shows the content of the current DHCPv6 snooping table.
DHCPv6 Snooping can be configured at Configuration→DHCPv6→Snooping.

DHCPv6 Snooping Table

This table display the currently known DHCPv6 clients and their assigned addresses.

Total entries: 0

Client DUID	MAC Address	Ingress Port	IAID	VLAN ID	Assigned Address	Lease Time	DHCP Server Address

Parameter	Description
Client DUID	The DHCP Unique Identifier (DUID) for the client. DHCPv6 uses this value to uniquely identify a client host instead of just using the MAC address of one of its interface ports (as DHCPv4 does).
MAC Address	The MAC address for the client interface port that sent the DHCPv6 message.
VLAN ID	The VLAN ID which is used by the client messages.
Local Ingress Port	The local port on the snooping switch where client messages are received.
DHCP Server Address	The IPv6 address of the DHCP server which assigned the address to the client.
IAID	Each client may contain multiple interfaces and may request addresses for each of these in the same DHCPv6 message. The Identity Association ID (IAID) value uniquely identifies the interface in the scope of the client.
Assigned Address	The address assigned to the interface identified by the IAID value.
Lease Time	The lease time associated with the assigned address in seconds.

5.11.2 Snooping Statistics

This shows statistics for DHCP snooping. It contains both RX and TX counters for all known DHCPv6 message types.
DHCPv6 Snooping can be configured at Configuration→DHCPv6→Snooping.

DHCPv6 Snooping Statistics

Selected port: Gi 1/1 Auto-refresh Refresh Clear

Receive Packets		Transmit Packets	
Rx Solicit	0	Tx Solicit	0
Rx Request	0	Tx Request	0
Rx InfoRequest	0	Tx InfoRequest	0
Rx Confirm	0	Tx Confirm	0
Rx Renew	0	Tx Renew	0
Rx Rebind	0	Tx Rebind	0
Rx Decline	0	Tx Decline	0
Rx Advertise	0	Tx Advertise	0
Rx Reply	0	Tx Reply	0
Rx Reconfigure	0	Tx Reconfigure	0
Rx Release	0	Tx Release	0
Rx DiscardUntrust	0		

Parameter	Description
Rx Solicit	Number of received DHCPv6 solicit messages. A client uses solicit to locate DHCPv6 servers that can provide configuration information.
Rx Request	Number of received DHCPv6 request messages. A client uses request to ask a server to assign or confirm IPv6 addresses and other configuration information offered previously.
Rx InfoRequest	Number of received DHCPv6 information-request messages. A client uses information-request to obtain configuration parameters (e.g., DNS information) without requesting IPv6 addresses.
Rx Confirm	Number of received DHCPv6 confirm messages. A client uses confirm to verify that previously assigned IPv6 addresses are still valid on the client's current link.
Rx Renew	Number of received DHCPv6 renew messages. A client uses renew to extend the lifetimes of existing IPv6 address bindings from the server that originally provided them.
Rx Rebind	Number of received DHCPv6 rebind messages. A client uses rebind to extend the lifetimes of existing address bindings when the original server is not reachable.
Rx Decline	Number of received DHCPv6 decline messages. A client uses decline to notify a server that one or more assigned IPv6 addresses are not appropriate for use on the link.
Rx Advertise	Number of received DHCPv6 advertise messages. A server sends advertise in response to a SOLICIT message to indicate that it is available and can provide configuration information.
Rx Reply	Number of received DHCPv6 reply messages. A server uses reply to respond to messages such as request, renew, rebind, confirm, release, decline, and information-request
Rx Reconfigure	Number of received DHCPv6 reconfigure messages. A server uses reconfigure to instruct clients to initiate a new request or information- request exchange to obtain updated configuration.
Rx Release	Number of received DHCPv6 release messages. A client uses release to inform a server that it will no longer use one or more IPv6 address bindings.
Rx DiscardUntrust	The *DiscardUntrust* counter indicate the number of received DHCP server packets that has been discarded due to the port being untrusted.
Tx Solicit	Number of transmitted DHCPv6 solicit messages observed by the switch (e.g., forwarded SOLICIT messages from clients).
Tx Request	Number of transmitted DHCPv6 request messages observed by the switch (e.g., forwarded REQUEST messages from clients).
Tx InfoRequest	Number of transmitted DHCPv6 information - request messages observed by the switch.
Tx Confirm	Number of transmitted DHCPv6 confirm messages observed by the switch.
Tx Renew	Number of transmitted DHCPv6 renew messages observed by the switch.
Tx Rebind	Number of transmitted DHCPv6 rebind messages observed by the switch.
Tx Decline	Number of transmitted DHCPv6 decline messages observed by the switch.
Tx Advertise	Number of transmitted DHCPv6 advertise messages observed by the switch (e.g.,forwarded advertise messages from servers).
Tx Reply	Number of transmitted DHCPv6 reply messages observed by the switch.
Tx Reconfigure	Number of transmitted DHCPv6 reconfigure messages observed by the switch.
Tx Release	Number of transmitted DHCPv6 release messages observed by the switch.

5.11.3 Relay

This shows current, configured relay agents and their statistics.
DHCPv6 relay can be configured at Configuration→DHCPv6→Relay.

DHCPv6 Relay Status and Statistics

Auto-refresh ☐ Refresh

Dropped server packets with interface option missing: **0**

Interface	Relay Interface	Relay Address	Tx to server	Rx from server	Server pkts dropped	Tx to client	Rx from client	Client pkts dropped	Clear stats
No entry exists									

Clear all statistics

Parameter	Description
Interface	Interface identification. The id of the interface that receives client requests.
Relay Interface	Interface identification. The id of the interface used for relaying.
Relay Address	An Ipv6 address represented as human readable test as specified in RFC5952. The IPv6 address that requests shall be relayed to. The default value 'ff05::1:3' means 'any DHCPv6 server'.
Tx to server	Integer number. Number of packets relayed to server.
Rx from server	Integer number. Number of packets received from server.
Server pkts dropped	Integer number. Number of packets from server that relay agent drops.
Tx to client	Integer number. Number of packets sent to client.
Rx from client	Integer number. Number of packets received from client.
Client pkts dropped	Integer number. Number of packets from client that relay agent drops.
Clear stats	Resets all statistics counters of relevant entry to zero.

5.12 Security

5.12.1 Switch

5.12.1.1 Access Management Statistics

This shows statistics for access management.

Access Management can be configured at Configuration→Switch→Access Management.

Access Management Statistics

Auto-refresh ☐ Refresh Clear

Interface	Received Packets	Allowed Packets	Discarded Packets
HTTP	0	0	0
HTTPS	0	0	0
SNMP	0	0	0
TELNET	0	0	0
SSH	0	0	0

Parameter	Description
Interface	The interface type through which the remote host can access the switch.
Received Packets	Number of received packets from the interface when access management mode is enabled.
Allowed Packets	Number of allowed packets from the interface when access management mode is enabled.

Discarded Packets	Number of discarded packets from the interface when access management mode is enabled.

5.12.1.2 RMON

Statistics

This shows an overview of RMON Statistics entries.
RMON Statistics can be configured at Configuration→Switch→RMON→Statistics.

RMON Statistics Status Overview

Auto-refresh ☐ Refresh |<< >>

Start from Control Index 0 with 20 entries per page.

ID	Data Source (ifIndex)	Drop	Octets	Pkts	Broad-cast	Multi-cast	CRC Errors	Under-size	Over-size	Frag.	Jabb.	Coll.	64 Bytes	65 ~ 127	128 ~ 255	256 ~ 511	512 ~ 1023	1024 ~ 1588
No more entries																		

Parameter	Description
ID	Indicates the index of Statistics entry.
Data Source(ifIndex)	The port ID which wants to be monitored.
Drop	The total number of events in which packets were dropped by the probe due to lack of resources.
Octets	The total number of octets of data (including those in bad packets) received on the network.
Pkts	The total number of packets (including bad packets, broadcast packets, and multicast packets) received.
Broad-cast	The total number of good packets received that were directed to the broadcast address.
Multi-cast	The total number of good packets received that were directed to a multicast address.
CRC 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 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).
Under-size	The total number of packets received that were less than 64 octets.
Over-size	The total number of packets received that were longer than 1518 octets.
Frag.	The number of frames which size is less than 64 octets received with invalid CRC.
Jabb.	The number of frames which size is larger than 64 octets received with invalid CRC.
Coll.	The best estimate of the total number of collisions on this Ethernet segment.
64	The total number of packets (including bad packets) received that were 64 octets in length.
65~127	The total number of packets (including bad packets) received that were between 65 to 127 octets in length.
128~255	The total number of packets (including bad packets) received that were between 128 to 255 octets in length.
256~511	The total number of packets (including bad packets) received that were between 256 to 511 octets in length.
512~1023	The total number of packets (including bad packets) received that were between 512 to 1023 octets in length.
1024~1588	The total number of packets (including bad packets) received that were between 1024 to 1588 octets in length.

History

This shows an overview of RMON History entries.
RMON History can be configured at Configuration→Switch→RMON→History.

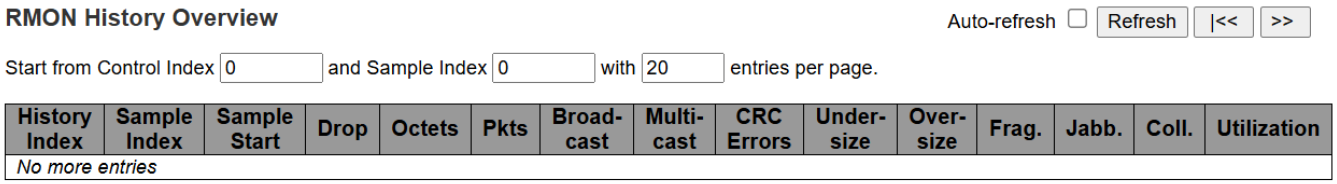

Parameter	Description
History Index	Indicates the index of History control entry.
Sample Index	Indicates the index of the data entry associated with the control entry.
Sample Start	The value of sysUpTime at the start of the interval over which this sample was measured.
Drop	The total number of events in which packets were dropped by the probe due to lack of resources.
Octets	The total number of octets of data (including those in bad packets) received on the network.
Pkts	The total number of packets (including bad packets, broadcast packets, and multicast packets) received.
Broadcast	The total number of good packets received that were directed to the broadcast address.
Multicast	The total number of good packets received that were directed to a multicast address.
CRCErrors	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 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).
Undersize	The total number of packets received that were less than 64 octets.
Oversize	The total number of packets received that were longer than 1518 octets.
Frag.	The number of frames which size is less than 64 octets received with invalid CRC.
Jabb.	The number of frames which size is larger than 64 octets received with invalid CRC.
Coll.	The best estimate of the total number of collisions on this Ethernet segment.
Utilization	The best estimate of the mean physical layer network utilization on this interface during this sampling interval, in hundredths of a percent.

Alarm

This shows an overview of RMON Alarm entries.
RMON Alarm can be configured at Configuration→Switch→RMON→Alarm.

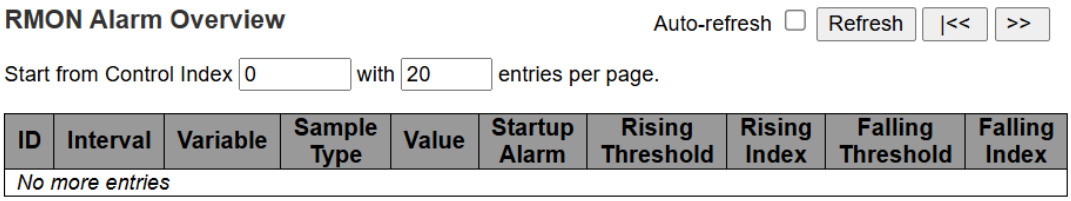

Parameter	Description
ID	Indicates the index of Alarm control entry.
Interval	Indicates the interval in seconds for sampling and comparing the rising and falling threshold.
Variable	Indicates the particular variable to be sampled.
Sample Type	The method of sampling the selected variable and calculating the value to be compared against the thresholds.
Value	The value of the statistic during the last sampling period.
Startup Alarm	The alarm that may be sent when this entry is first set to valid.
Rising Threshold	Rising threshold value.
Rising Index	Rising event index.
Falling Threshold	Falling threshold value.
Falling Index	Falling event index.

Event

This shows an overview of RMON Event table entries.
RMON Event can be configured at Configuration→Switch→RMON→Event.

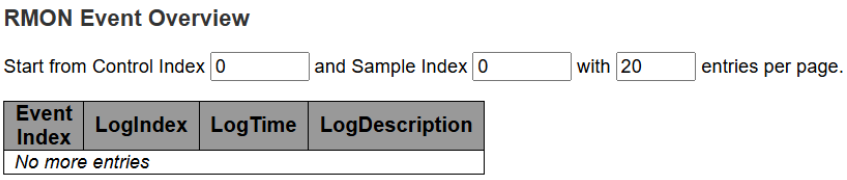

Parameter	Description
Event Index	Indicates the index of the event entry.
Log Index	Indicates the index of the log entry.
LogTime	Indicates Event log time
LogDescription	Indicates the Event description.

5.12.2 Network

5.12.2.1 Port Security

This shows the Port Security status Port Security may be configured both administratively and indirectly through other software modules - the so-called user modules. When a user module has enabled port security on a port, the port is set-up for software-based learning. In this mode, frames from unknown MAC addresses are passed on to the port security module, which in turn asks all user modules whether to allow this new MAC address to forward or block it. For a MAC address to be set in the forwarding state, all enabled user modules must unanimously agree on allowing the MAC address to forward. If only one chooses to block it, it will be blocked until that user module decides otherwise.

5.12.2.2 Overview

The status page is divided into two sections - one with a legend of user modules and one with the actual port status.
Port Security can be configured at Configuration→Network→Port Security.

Port Security Switch Status

User Module Legend

User Module Name	Abbr
Port Security (Admin)	P
802.1X	8
Voice VLAN	V

Port Status

Clear	Port	Users	Violation Mode	State	MAC Count		
					Current	Violating	Limit
Clear	1	---	Disabled	Disabled	-	-	-
Clear	2	---	Disabled	Disabled	-	-	-
Clear	3	---	Disabled	Disabled	-	-	-
Clear	4	---	Disabled	Disabled	-	-	-
Clear	5	---	Disabled	Disabled	-	-	-
Clear	6	---	Disabled	Disabled	-	-	-
Clear	7	---	Disabled	Disabled	-	-	-
Clear	8	---	Disabled	Disabled	-	-	-
Clear	9	---	Disabled	Disabled	-	-	-
Clear	10	---	Disabled	Disabled	-	-	-
Clear	11	---	Disabled	Disabled	-	-	-
Clear	12	---	Disabled	Disabled	-	-	-

Parameter	Description
User Module Name	The full name of a module that may request Port Security services.
Abbr	A one-letter abbreviation of the user module. This is used in the Users column in the port status table.
Clear	Click to remove all dynamic MAC addresses on all VLANs on this port. The button is only clickable if number of secured MAC addresses is non-zero.
Users	Each of the user modules has a column that shows whether that module has enabled Port Security or not. A '-' means that the corresponding user module is not enabled, whereas a letter indicates that the user module abbreviated by that letter (see Abbr) has enabled port security.
Violation Mode	Shows the configured Violation Mode of the port. It can take one of four values: **Disabled**: Port Security is not administratively enabled on this port. **Protect**: Port Security is administratively enabled in Protect mode. **Restrict**: Port Security is administratively enabled in Restrict mode. **Shutdown**: Port Security is administratively enabled in Shutdown mode.

State	Shows the current state of the port. It can take one of four values: **Disabled**: No user modules are currently using the Port Security service. **Ready**: The Port Security service is in use by at least one user module, and is awaiting frames from unknown MAC addresses to arrive. **Limit Reached**: The Port Security service is administratively enabled and the limit is reached. **Shut down**: The Port Security service is administratively enabled and the port is shut down. No MAC addresses can be learned on the port until it is administratively re-opened by administratively taking the port down and then back up on the Configuration→Ports page. Alternatively, the switch may be booted or reconfigured Port Security-wise.
MAC Count (Current, Violating, Limit)	The three columns indicate the number of currently learned MAC addresses (forwarding as well as blocked), the number of violating MAC address (only counting in Restrict mode) and the maximum number of MAC addresses that can be learned on the port, respectively. If no user modules are enabled on the port, the Current column will show a dash (-). If Port Security is not administratively enabled on the port, the Violating and Limit columns will show a dash (-).

5.12.2.3 Details

This shows the MAC address secured by the Port Security module. Port Security Port Security may be configured both administratively and indirectly through other software modules - the so-called user modules. When a user module has enabled port security on a port, the port is set-up for software-based learning. In this mode, frames from unknown MAC addresses are passed on to the Port Security module, which in turn asks all user modules whether to allow this new MAC address to forward or block it. For a MAC address to be set in the forwarding state, all enabled user modules must unanimously agree on allowing the MAC address to forward. If only one chooses to block it, it will be blocked until that user module decides otherwise.

Notice that if you have added static or sticky MAC addresses, they will show up on this page only if Port Security is enabled on the interface to which they pertain.
Port Security can be configured at Configuration→Network→Port Security.

Port Security Port Status All Ports

Delete	Port	VLAN ID	MAC Address	Type	State	Age/Hold
No MAC addresses attached						

Parameter	Description
Delete	Click to remove this particular MAC addresses from MAC address table. The button is only clickable if the entry type is Dynamic. Use the "Configuration→Security→Port Security→MAC Addresses" page to remove Static and Sticky entries.
Port	If all ports are shown (can be selected through the drop-down box on the top right), this one shows the port to which the MAC address is bound.
VLAN ID & MAC Address	The VLAN ID and MAC address that is seen on this port. If no MAC addresses are learned, a single row stating *"No MAC addresses attached"* is displayed.
Type	Indicates the type of entry. Takes one of three values: **Dynamic**: The entry is learned through learn frames coming to the Port Security module while the port in question is not in sticky mode. **Static**: The entry is entered by the end-user through management. Entry is not subject to aging. **Sticky**: When the port is in sticky mode, all entries that would otherwise have been learned as dynamic are learned as sticky. Sticky entries are part of the running-config and can therefore be saved to startup-config. An important aspect of sticky MAC addresses is that they survive link changes (in contrast to Dynamic, which will have to be learned again). They also survive reboots if running-config is saved to startup-config.

State	Indicates whether the corresponding MAC address is violating (administrative user has configured the interface in "Restrict" mode and the MAC address is blocked), blocked, or forwarding.
Age/Hold	If at least one user module has decided to block this MAC address, it will stay in the blocked state until the hold time (measured in seconds) expires. If all user modules have decided to allow this MAC address to forward, and aging is enabled, the Port Security module will periodically check that this MAC address still forwards traffic. If the age period (measured in seconds) expires and no frames have been seen, the MAC address will be removed from the MAC address table. Otherwise, a new age period will begin. If aging is disabled or a user module has decided to hold the MAC address indefinitely, a dash (-) will be shown.

5.12.2.4 NAS

NAS is an acronym for **N**etwork **A**ccess **S**erver. The NAS is meant to act as a gateway to guard access to a protected source. A client connects to the NAS, and the NAS connects to another resource asking whether the client's supplied credentials are valid. Based on the answer, the NAS then allows or disallows access to the protected resource. An example of a NAS implementation is IEEE 802.1X.

5.12.2.5 Switch

This shows an overview of the current NAS port states.
NAS switch can be configured at Configuration→Network→NAS.

Network Access Server Switch Status

Port	Admin State	Port State	Last Source	Last ID	QoS Class	Port VLAN ID
1	Force Authorized	Globally Disabled			-	
2	Force Authorized	Globally Disabled			-	
3	Force Authorized	Globally Disabled			-	
4	Force Authorized	Globally Disabled			-	
5	Force Authorized	Globally Disabled			-	
6	Force Authorized	Globally Disabled			-	
7	Force Authorized	Globally Disabled			-	
8	Force Authorized	Globally Disabled			-	
9	Force Authorized	Globally Disabled			-	
10	Force Authorized	Globally Disabled			-	
11	Force Authorized	Globally Disabled			-	
12	Force Authorized	Globally Disabled			-	

Parameter	Description
Port	The switch port number. Click to navigate to detailed NAS statistics for this port.
Admin State	The port's current administrative state. Refer to NAS Admin State for a description of possible values.
Port State	The current state of the port. Refer to NAS Port State for a description of the individual states.
Last Source	The source MAC address carried in the most recently received EAPOL frame for EAPOL-based authentication, and the most recently received frame from a new client for MAC-based authentication.
Last ID	The user name (supplicant identity) carried in the most recently received Response Identity EAPOL frame for EAPOL-based authentication, and the source MAC address from the most recently received frame from a new client for MAC-based authentication.
QoS Class	QoS Class assigned to the port by the RADIUS server if enabled.
Port VLAN ID	The VLAN ID that NAS has put the port in. The field is blank, if the Port VLAN ID is not overridden by NAS. If the VLAN ID is assigned by the RADIUS server, "(RADIUS-assigned)" is appended to the VLAN ID. Read more about RADIUS-assigned VLANs here. If the port is moved to the Guest VLAN, "(Guest)" is appended to the VLAN ID. Read more about Guest VLANs here.

5.12.2.6 Port

This page provides detailed NAS statistics for a specific switch port running EAPOL-based IEEE 802.1X authentication. For MAC-based ports, it shows selected backend server (RADIUS Authentication Server) statistics, only. Use the port select box to select which port details to be displayed
NAS Port can be configured at Configuration→Network→NAS.

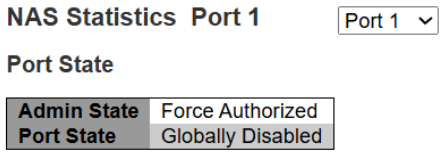

Parameter	Description
Admin State	The port's current administrative state. Refer to NAS Admin State for a description of possible values.
Port State	The current state of the port. Refer to NAS Port State for a description of the individual states.
QoS Class	The QoS class assigned by the RADIUS server. The field is blank if no QoS class is assigned.
Port VLAN ID	The VLAN ID that NAS has put the port in. The field is blank, if the Port VLAN ID is not overridden by NAS. If the VLAN ID is assigned by the RADIUS server, "(RADIUS-assigned)" is appended to the VLAN ID. Read more about RADIUS-assigned VLANs here. If the port is moved to the Guest VLAN, "(Guest)" is appended to the VLAN ID. Read more about Guest VLANs here.

EAPOL Counters

These supplicant frame counters are available for the following administrative states:

- Force Authorized
- Force Unauthorized
- Port-based 802.1X
- Single 802.1X
- Multi 802.1X

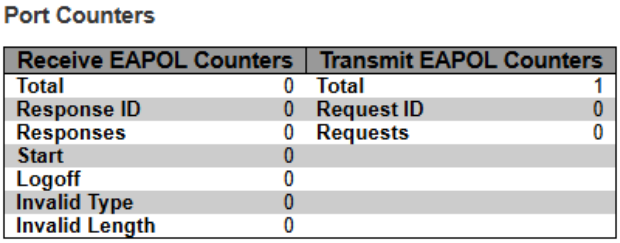

Direction	Name	IEEE Name	Description
Rx	**Total**	dot1xAuthEapolFramesRx	The number of valid EAPOL frames of any type that have been received by the switch.
Rx	**Response ID**	dot1xAuthEapolRespIdFramesRx	The number of valid EAPOL Response Identity frames that have been received by the switch.
Rx	**Responses**	dot1xAuthEapolRespFramesRx	The number of valid EAPOL response frames (other than Response Identity frames) that have been received by the switch.
Rx	**Start**	dot1xAuthEapolLogoffFramesRx	The number of EAPOL Start frames that have been received by the switch.
Rx	**Logoff**	dot1xAuthEapolLogoffFramesRx	The number of valid EAPOL Logoff frames that have been received by the switch.
Rx	**Invalid Type**	dot1xAuthInvalidEapolFramesRx	The number of EAPOL frames that have been received by the switch in which the frame type is not recognized.

Rx	**Invalid Length**	dot1xAuthEapLengthErrorFramesRx	The number of EAPOL frames that have been received by the switch in which the Packet Body Length field is invalid.
Tx	**Total**	dot1xAuthEapolFramesTx	The number of EAPOL frames of any type that have been transmitted by the switch.
Tx	**Request ID**	dot1xAuthEapolReqIdFramesTx	The number of EAPOL Request Identity frames that have been transmitted by the switch.
Tx	**Requests**	dot1xAuthEapolReqFramesTx	The number of valid EAPOL Request frames (other than Request Identity frames) that have been transmitted by the switch.

Backend Server Counters

These backend (RADIUS) frame counters are available for the following administrative states:

- Port-based 802.1X
- Single 802.1X
- Multi 802.1X
- MAC-based Auth.

Receive Backend Server Counters		Transmit Backend Server Counters	
Access Challenges	0	Responses	0
Other Requests	1		
Auth. Successes	0		
Auth. Failures	0		

Direction	Name	IEEE Name	Description
Rx	**Access Challenges**	dot1xAuthBackendAccessChallenges	**802.1X-based:** Counts the number of times that the switch receives the first request from the backend server following the first response from the supplicant. Indicates that the backend server has communication with the switch. **MAC-based:** Counts all Access Challenges received from the backend server for this port (left-most table) or client (right-most table).
Rx	**Other Requests**	dot1xAuthBackendOtherRequestsToSupplicant	**802.1X-based:** Counts the number of times that the switch sends an EAP Request packet following the first to the supplicant. Indicates that the backend server chose an EAP-method. **MAC-based:** Not applicable.
Rx	**Auth. Successes**	dot1xAuthBackendAuthSuccesses	**802.1X- and MAC-based:** Counts the number of times that the switch receives a success indication. Indicates that the supplicant or client has successfully authenticated to the backend server.
Rx	**Auth. Failures**	dot1xAuthBackendAuthFails	**802.1X- and MAC-based:** Counts the number of times that the switch receives a failure message. This indicates that the supplicant or client has not authenticated to the backend server.
Tx	**Responses**	dot1xAuthBackendResponses	**802.1X-based:** Counts the number of times that the switch attempts to send a supplicant's first response packet to the backend server. Indicates the switch attempted communication with the backend server. Possible retransmissions are not counted. **MAC-based:** Counts all the backend server packets sent from

			the switch towards the backend server for a given port (left-most table) or client (right-most table). Possible retransmissions are not counted.

Last Supplicant/Client Info

These backend (RADIUS) frame counters are available for the following administrative states:

- Port-based 802.1X
- Single 802.1X
- Multi 802.1X
- MAC-based Auth.

Last Supplicant Info	
MAC Address	00-00-00-00-00-00
VLAN ID	0
Version	0
Identity	

Name	IEEE Name	Description
MAC Address	dot1xAuthLastEapolFrameSource	The MAC address of the last supplicant/client.
VLAN ID		The VLAN ID on which the last frame from the last supplicant/client was received.
Version	dot1xAuthLastEapolFrameVersion	**802.1X-based:** The protocol version number carried in the most recently received EAPOL frame. **MAC-based:** Not applicable.
Identity		**802.1X-based:** The user name (supplicant identity) carried in the most recently received Response Identity EAPOL frame. **MAC-based:** Not applicable.

Selected Counters

The Selected Counters table is visible when the port is in one of the following administrative states:

- Multi 802.1X
- MAC-based Auth.

Selected Counters

Receive Backend Server Counters	Transmit Backend Server Counters
Access Challenges	Responses
Auth. Successes	
Auth. Failures	
Client Info	
MAC Address	*No client selected*
VLAN ID	

The table is identical to and is placed next to the Port Counters table, and will be empty if no MAC address is currently selected. To populate the table, select one of the attached MAC Addresses from the table below.

Attached MAC Address

Attached Clients

MAC Address	VLAN ID	State	Last Authentication
f0-12-04-14-d2-36	1	Unauthorized	

Parameter	Description

Identity	Shows the identity of the supplicant, as received in the Response Identity EAPOL frame. Clicking the link causes the supplicant's EAPOL and Backend Server counters to be shown in the Selected Counters table. If no supplicants are attached, it shows *No supplicants attached*. This column is not available for MAC-based Auth.
MAC Address	For Multi 802.1X, this column holds the MAC address of the attached supplicant. For MAC-based Auth., this column holds the MAC address of the attached client. Clicking the link causes the client's Backend Server counters to be shown in the Selected Counters table. If no clients are attached, it shows *No clients attached*.
VLAN ID	This column holds the VLAN ID that the corresponding client is currently secured through the Port Security module.
State	The client can either be authenticated or unauthenticated. In the authenticated state, it is allowed to forward frames on the port, and in the unauthenticated state, it is blocked. As long as the backend server hasn't successfully authenticated the client, it is unauthenticated. If an authentication fails for one or the other reason, the client will remain in the unauthenticated state for Hold time seconds.
Last Authentication	Shows the date and time of the last authentication of the client (successful as well as unsuccessful).

5.12.2.7 ACL Status

This shows the ACL status by different ACL users. Each row describes the ACE that is defined. It is a conflict if a specific ACE is not applied to the hardware due to hardware limitations. The maximum number of ACEs is **512** on each switch.
ACL can be configured at Configuration→Security→Network→ACL.

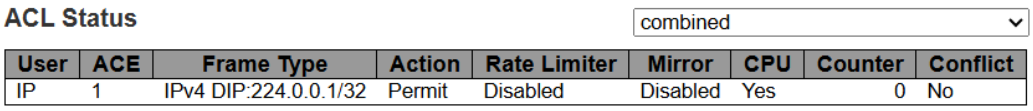

Parameter	Description
User	Indicates the ACL user.
ACE	Indicates the ACE ID on local switch.
Frame Type	Indicates the frame type of the ACE. Possible values are: **Any**: The ACE will match any frame type. **EType**: The ACE will match Ethernet Type frames. Note that an Ethernet Type based ACE will not get matched by IP and ARP frames. **ARP**: The ACE will match ARP/RARP frames. **IPv4**: The ACE will match all IPv4 frames. **IPv4/ICMP**: The ACE will match IPv4 frames with ICMP protocol. **IPv4/UDP:** The ACE will match IPv4 frames with UDP protocol. **IPv4/TCP**: The ACE will match IPv4 frames with TCP protocol. **IPv4/Other**: The ACE will match IPv4 frames, which are not ICMP/UDP/TCP. **IPv6**: The ACE will match all IPv6 standard frames
Action	Indicates the forwarding action of the ACE. **Permit**: Frames matching the ACE may be forwarded and learned. **Deny**: Frames matching the ACE are dropped. **Filter**: Frames matching the ACE are filtered.
Rate Limiter	Indicates the rate limiter number of the ACE. The allowed range is **1** to **15**. When **Disabled** is displayed, the rate limiter operation is disabled.
CPU	Forward packet that matched the specific ACE to CPU.
Counter	The counter indicates the number of times the ACE was hit by a frame.
Conflict	Indicates the hardware status of the specific ACE. The specific ACE is not applied to the hardware due to hardware limitations.

5.12.2.8 IP Source Guard

This shows the Dynamic IPv4 Source Guard Table entries. The Dynamic IP Source Guard Table is sorted first by port, then by VLAN ID, then by IP address, and then by MAC address. The IPv6 Source Guard uses identical parameters.

IP Source Guard can be configured at Configuration→Security→Network→IP Source Guard.

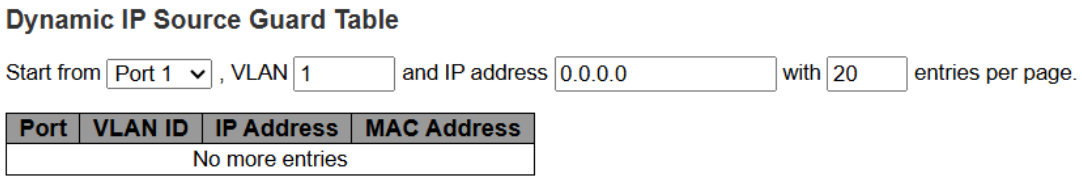

Parameter	Description
Port	Switch Port Number for which the entries are displayed.
VLAN ID	VLAN-ID in which the IP traffic is permitted.
IP Address	User IP address of the entry.
MAC Address	Source MAC address.

5.12.2.9 IPv6 Source Guard

The IPv6 Source Guard uses identical parameters to the IPv4 Source Guard.
For a detailed description, please refer to Monitor→Security→Network→IP Source Guard.

5.12.2.10 ARP Inspection

This shows the Dynamic IPv6 Source Guard Table entries. The Dynamic ARP Inspection Table contains up to 256 entries, and is sorted first by port, then by VLAN ID, then by MAC address, and the by IP address. All dynamic entries are learned from DHCP Snooping.
ARP Inspection can be configured at Configuration→Security→Network→ARP Inspection.

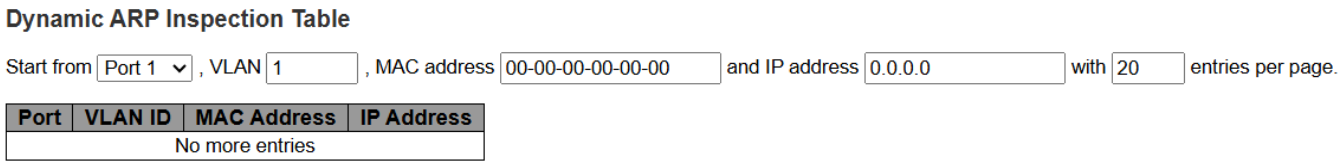

Parameter	Description
Port	Switch Port Number for which the entries are displayed.
VLAN ID	VLAN-ID in which the ARP traffic is permitted.
MAC Address	User MAC address of the entry.
IP Address	User IP address of the entry.

5.12.3 AAA

5.12.3.1 RADIUS Overview

This shows an overview of the status of the RADIUS servers configurable on the Authentication configuration page. RADIUS can be configured at Configuration→Security→AAA→RADIUS.

RADIUS Server Status Overview

#	IP Address	Authentication Port	Authentication Status	Accounting Port	Accounting Status
1			Disabled		Disabled
2			Disabled		Disabled
3			Disabled		Disabled
4			Disabled		Disabled
5			Disabled		Disabled

Parameter	Description
#	The RADIUS server number. Click to navigate to detailed statistics for this server.
IP Address	The IP address of this server.
Authentication Port	UDP port number for authentication.
Authentication Status	The current status of the server. This field takes one of the following values: **Disabled**: The server is disabled. **Not Ready:** The server is enabled, but IP communication is not yet up and running. **Ready**: The server is enabled, IP communication is up and running, and the RADIUS module is ready to accept access attempts. **Dead (X seconds left)**: Access attempts were made to this server, but it did not reply within the configured timeout. The server has temporarily been disabled, but will get re-enabled when the dead-time expires. The number of seconds left before this occurs is displayed in parentheses. This state is only reachable when more than one server is enabled.
Accounting Port	UDP port number for accounting.
Accounting Status	The current status of the server. This field takes one of the following values: **Disabled**: The server is disabled. **Not Ready:** The server is enabled, but IP communication is not yet up and running. **Ready**: The server is enabled, IP communication is up and running, and the RADIUS module is ready to accept access attempts. **Dead (X seconds left)**: Access attempts were made to this server, but it did not reply within the configured timeout. The server has temporarily been disabled, but will get re-enabled when the dead-time expires. The number of seconds left before this occurs is displayed in parentheses. This state is only reachable when more than one server is enabled.

5.12.3.2 RADIUS Details

This shows detailed statistics for a particular RADIUS server.
RADIUS can be configured at Configuration→Security→AAA→RADIUS.

RADIUS Authentication Statistics:
The statistics map closely to those specified in RFC4668 – RADIUS Authentication Client MIB.
Use the server select box to switch between the backend servers to show details for.

Packet Counters

RADIUS authentication server packet counter. There are seven receive and four transmit counters.

Receive Packets		Transmit Packets	
Access Accepts	0	Access Requests	0
Access Rejects	0	Access Retransmissions	0
Access Challenges	0	Pending Requests	0
Malformed Access Responses	0	Timeouts	0
Bad Authenticators	0		
Unknown Types	0		
Packets Dropped	0		

Direction	Name	RFC4668 Name	Description
Rx	**Access Accepts**	radiusAuthClientExtAccessAccepts	The number of RADIUS Access-Accept packets (valid or invalid) received from the server.
Rx	**Access Rejects**	radiusAuthClientExtAccessRejects	The number of RADIUS Access-Reject packets (valid or invalid) received from the server.
Rx	**Access Challenges**	radiusAuthClientExtAccessChallenges	The number of RADIUS Access-Challenge packets (valid or invalid) received from the server.
Rx	**Malformed Access Responses**	radiusAuthClientExtMalformedAccessResponses	The number of malformed RADIUS Access-Response packets received from the server. Malformed packets include packets with an invalid length. Bad authenticators or Message Authenticator attributes or unknown types are not included as malformed access responses.
Rx	**Bad Authenticators**	radiusAuthClientExtBadAuthenticators	The number of RADIUS Access-Response packets containing invalid authenticators or Message Authenticator attributes received from the server.
Rx	**Unknown Types**	radiusAuthClientExtUnknownTypes	The number of RADIUS packets that were received with unknown types from the server on the authentication port and dropped.
Rx	**Packets Dropped**	radiusAuthClientExtPacketsDropped	The number of RADIUS packets that were received from the server on the authentication port and dropped for some other reason.
Tx	**Access Requests**	radiusAuthClientExtAccessRequests	The number of RADIUS Access-Request packets sent to the server. This does not include retransmissions.
Tx	**Access Retransmissions**	radiusAuthClientExtAccessRetransmissions	The number of RADIUS Access-Request packets retransmitted to the RADIUS authentication server.
Tx	**Pending Requests**	radiusAuthClientExtPendingRequests	The number of RADIUS Access-Request packets destined for the server that have not yet timed out or received a response. This variable is incremented when an Access-Request is sent and decremented due to receipt of an

			Access-Accept, Access-Reject, Access-Challenge, timeout, or retransmission.
Tx	**Timeouts**	radiusAuthClientExtTimeouts	The number of authentication timeouts to the server. After a timeout, the client may retry to the same server, send to a different server, or give up. A retry to the same server is counted as a retransmit as well as a timeout. A send to a different server is counted as a Request as well as a timeout.

Other Info

This section contains information about the state of the server and the latest round-trip time.

Other Info	
IP Address	
State	Disabled
Round-Trip Time	0 ms

Name	RFC4668 Name	Description
IP Address		IP address and UDP port for the authentication server in question.
State		Shows the state of the server. It takes one of the following values: **Disabled**: The selected server is disabled. **Not Ready**: The server is enabled, but IP communication is not yet up and running. **Ready**: The server is enabled, IP communication is up and running, and the RADIUS module is ready to accept access attempts. **Dead (X seconds left)**: Access attempts were made to this server, but it did not reply within the configured timeout. The server has temporarily been disabled, but will get re-enabled when the dead-time expires. The number of seconds left before this occurs is displayed in parentheses. This state is only reachable when more than one server is enabled.
Round-Trip Time	radiusAuthClientExtRoundTripTime	The time interval (measured in milliseconds) between the most recent Access-Reply/Access-Challenge and the Access-Request that matched it from the RADIUS authentication server. The granularity of this measurement is 100 ms. A value of 0 ms indicates that there hasn't been round-trip communication with the server yet.

RADIUS Accounting Statistics:
The statistics map closely to those specified in RFC4670 – RADIUS Accounting Client MIB.
Use the server select box to switch between the backend servers to show details for.

Packet Counters

RADIUS accounting server packet counter. There are five receive and four transmit counters.

Receive Packets		Transmit Packets	
Responses	0	**Requests**	0
Malformed Responses	0	**Retransmissions**	0
Bad Authenticators	0	**Pending Requests**	0
Unknown Types	0	**Timeouts**	0
Packets Dropped	0		

Direction	Name	RFC4668 Name	Description
Rx	**Responses**	radiusAccClientExtResponses	The number of RADIUS packets (valid or invalid) received from the server.

Rx	**Malformed Responses**	radiusAccClientExtMalformedResponses	The number of malformed RADIUS packets received from the server. Malformed packets include packets with an invalid length. Bad authenticators or unknown types are not included as malformed access responses.
Rx	**Bad Authenticators**	radiusAcctClientExtBadAuthenticators	The number of RADIUS packets containing invalid authenticators received from the server.
Rx	**Unknown Types**	radiusAccClientExtUnknownTypes	The number of RADIUS packets of unknown types that were received from the server on the accounting port.
Rx	**Packets Dropped**	radiusAccClientExtPacketsDropped	The number of RADIUS packets that were received from the server on the accounting port and dropped for some other reason.
Tx	**Requests**	radiusAccClientExtRequests	The number of RADIUS packets sent to the server. This does not include retransmissions.
Tx	**Retransmissions**	radiusAccClientExtRetransmissions	The number of RADIUS packets retransmitted to the RADIUS accounting server.
Tx	**Pending Requests**	radiusAccClientExtPendingRequests	The number of RADIUS packets destined for the server that have not yet timed out or received a response. This variable is incremented when a Request is sent and decremented due to receipt of a Response, timeout, or retransmission.
Tx	**Timeouts**	radiusAccClientExtTimeouts	The number of accounting timeouts to the server. After a timeout, the client may retry to the same server, send to a different server, or give up. A retry to the same server is counted as a retransmit as well as a timeout. A send to a different server is counted as a Request as well as a timeout.

Other Info

This section contains information about the state of the server and the latest round-trip time.

Other Info	
IP Address	
State	Disabled
Round-Trip Time	0 ms

Name	RFC4670 Name	Description
IP Address		IP address and UDP port for the accounting server in question.

State		Shows the state of the server. It takes one of the following values: **Disabled**: The selected server is disabled. **Not Ready**: The server is enabled, but IP communication is not yet up and running. **Ready**: The server is enabled, IP communication is up and running, and the RADIUS module is ready to accept access attempts. **Dead (X seconds left)**: Access attempts were made to this server, but it did not reply within the configured timeout. The server has temporarily been disabled, but will get re-enabled when the dead-time expires. The number of seconds left before this occurs is displayed in parentheses. This state is only reachable when more than one server is enabled.
Round-Trip Time	radiusAuthClientExtRoundTripTime	The time interval (measured in milliseconds) between the most recent Response and the Request that matched it from the RADIUS accounting server. The granularity of this measurement is 100 ms. A value of 0 ms indicates that there hasn't been round-trip communication with the server yet.

5.13 Aggregation

Link Aggregation combines multiple ports in parallel to Increase bandwidth beyond the capacity of a single port and to provide redundancy for higher availability.

5.13.1 Status

This shows the status of ports in the Aggregation group.
Aggregation can be configured at Configuration→Aggregation.

Aggregation Status

Aggr ID	Name	Type	Speed	Configured Ports	Aggregated Ports
No aggregation groups					

Parameter	Description
Aggr ID	The Aggregation ID associated with this aggregation instance.
Name	Name of the Aggregation group ID.
Type	Type of the Aggregation group (Static or LACP).
Speed	Speed of the Aggregation group.
Configured ports	Configured member ports of the Aggregation group.
Aggregated ports	Aggregated member ports of the Aggregation group.

5.13.2 LACP

LACP is an IEEE 802.3ad standard protocol. The Link Aggregation Control Protocol, allows bundling several Physical ports together to form a single logical port.

5.13.2.1 System Status

This shows a status overview for the system-level LACP information.
LACP can be configured at Configuration→Aggregation→LACP.

LACP System Status

Local System ID

Priority	MAC Address
32768	00-0b-04-14-d8-bf

Partner System Status

Aggr ID	Partner System ID	Partner Prio	Partner Key	Last Changed	Local Ports
No ports enabled or no existing partners					

Parameter	Description
Priority	The LACP system priority value assigned to this switch.
MAC Address	The MAC address used as the local system ID in LACP.
Aggr ID	The Aggregation ID associated with this aggregation instance.
Partner System ID	The system ID (MAC address) of the aggregation partner.
Partner Prio	The priority that the partner has assigned to this aggregation ID.
Partner Key	The Key that the partner has assigned to this aggregation ID.
Last Changed	The time since this aggregation changed.
Local Ports	Shows which ports are a part of this aggregation for this switch.

5.13.2.2 Internal Status

This shows a status overview for the LACP internal (i.e. local system) status for all ports. Only ports that are part of an LACP group are shown. For details on the shown parameters please refer to IEEE 801.AX-2014.
LACP can be configured at Configuration→Aggregation→LACP.

LACP Internal Port Status

Port	State	Key	Priority	Activity	Timeout	Aggregation	Synchronization	Collecting	Distributing	Defaulted	Expired
No LACP ports enabled											

Parameter	Description
Port	The switch port number.
State	The current port state: • **Down:** The port is not active. • **Active:** The port is in active state. • **Standby:** The port is in standby state.
Key	The key assigned to this port. Only ports with the same key can aggregate together.
Priority	The priority assigned to this aggregation group.

Activity	The LACP mode of the group (Active or Passive).
Timeout	The timeout mode configured for the port (Fast or Slow).
Aggregation	Show whether the system considers this link to be "aggregateable"; i.e., a potential candidate for aggregation.
Synchronization	Show whether the system considers this link to be "IN_SYNC"; i.e., it has been allocated to the correct LAG, the group has been associated with a compatible Aggregator, and the identity of the LAG is consistent with the System ID and operational Key information transmitted.
Collecting	Show if collection of incoming frames on this link is enabled.
Distributing	Show if distribution of outgoing frames on this link is enabled.
Defaulted	Show if the Actor's Receive machine is using Defaulted operational Partner information.
Expired	Show if that the Actor's Receive machine is in the EXPIRED state.

5.13.2.3 Neighbor Status

This shows a status overview for the LACP neighbor status for all ports. Only ports that are part of an LACP group are shown. For details on the shown parameters please refer to IEEE 801.AX-2014.
LACP can be configured at Configuration→Aggregation→LACP.

LACP Neighbor Port Status

Port	State	Aggr ID	Partner Key	Partner Port	Partner Port Prio	Activity	Timeout	Aggregation	Synchronization	Collecting	Distributing	Defaulted	Expired
No LACP neighbor status available													

Parameter	Description
Port	The switch port number.
State	The current port state: • **Down:** The port is not active. • **Active:** The port is in active state. • **Standby:** The port is in standby state.
Aggr ID	The aggregation group ID which the port is assigned to.
Partner Key	The key assigned to this port by the partner.
Partner Port	The partner port number associated with this link.
Partner Port Priority	The priority assigned to this partner port.
Activity	The LACP mode of the group (Active or Passive).
Timeout	The timeout mode configured for the partner port (Fast or Slow).
Aggregation	Show whether the partner considers this link to be "aggregateable"; i.e., a potential candidate for aggregation.
Synchronization	Show whether the partner considers this link to be "IN_SYNC"; i.e., it has been allocated to the correct LAG, the group has been associated with a compatible Aggregator, and the identity of the LAG is consistent with the System ID and operational Key information transmitted.
Collecting	Show if collection of incoming frames on this link is enabled.
Distributing	Show if distribution of outgoing frames on this link is enabled.

Defaulted	Show if the partners Receive machine is using Defaulted operational Partner information.
Expired	Show if that the partners Receive machine is in the EXPIRED state.

5.13.2.4 Port Statistics

This shows a status overview for the LACP statistics for all ports.
LACP can be configured at Configuration→Aggregation→LACP.

LACP Statistics

Port	LACP Received	LACP Transmitted	Discarded	
			Unknown	Illegal
No ports enabled				

Parameter	Description
Port	The switch port number.
LACP Received	Shows how many LACP frames have been received at each port.
LACP Transmitted	Shows how many LACP frames have been sent from each port.
Discarded	Shows how many unknown or illegal LACP frames have been discarded at each port.

5.14 Loop Protection

This shows the loop protection status of all ports on the switch.
Loop Protection can be configured at Configuration→Loop Protection.

Loop Protection Status

Port	Action	Transmit	Loops	Status	Loop	Time of Last Loop
No ports enabled						

Parameter	Description
Port	The switch port number of the logical port.
Action	The currently configured port action.
Transmit	The currently configured port transmit mode.
Loops	The number of loops detected on this port.
Status	The current loop protection status of the port.
Loop	Whether a loop is currently detected on the port.
Time of Last Loop	The time of the last loop event detected.

5.15 Spanning Tree

Spanning Tree Protocol is an OSI layer-2 protocol which ensures a loop free topology for any bridged LAN.

5.15.1 Bridge Status

This shows a status overview of all STP bridge instances. The displayed table contains a row for each STP bridge instance, where the column displays the following information.
Spanning Tree can be configured at Configuration→Spanning Tree.

STP Bridges

MSTI	Bridge ID	Root			Topology Flag	Topology Change Last
		ID	Port	Cost		
CIST	32768.00-0B-04-14-D8-BF	32767.00-0B-04-14-1D-C6	8	20000	Steady	1d 04:50:01

Parameter	Description
MSTI	The Bridge Instance. This is also a link to the STP Detailed Bridge Status.
Bridge ID	The Bridge ID of this Bridge instance.
Root ID	The Bridge ID of the currently elected root bridge.
Root Port	The switch port currently assigned the *root* port role.
Root Cost	Root Path Cost. For the Root Bridge it is zero. For all other Bridges, it is the sum of the Port Path Costs on the least cost path to the Root Bridge.
Topology Flag	The current state of the Topology Change Flag of this Bridge instance.
Topology Change Last	The time since last Topology Change occurred.

STP Detailed Bridge Status

This show detailed information on a single STP bridge instance, along with port state for all active ports Associated.

STP Detailed Bridge Status

STP Bridge Status	
Bridge Instance	CIST
Bridge ID	32768.00-0B-04-14-D8-BF
Root ID	32767.00-0B-04-14-1D-C6
Root Cost	20000
Root Port	8
Regional Root	32768.00-0B-04-14-D8-BF
Internal Root Cost	0
Topology Flag	Steady
Topology Change Count	4
Topology Change Last	1d 04:52:49

CIST Ports & Aggregations State

Port	Port ID	Role	State	Path Cost	Edge	Point-to-Point	Uptime
8	128:008	RootPort	Forwarding	20000	No	Yes	1d 04:52:53

Parameter	Description
Bridge Instance	The Bridge instance – **CIST**, **MSTI**, …

Bridge ID	The Bridge ID of this Bridge instance.
Root ID	The Bridge ID of the currently elected root bridge.
Root Port	The switch port currently assigned the *root* port role.
Root Cost	Root Path Cost. For the Root Bridge this is zero. For all other Bridges, it is the sum of the Port Path Costs on the least cost path to the Root Bridge.
Regional Root	The Bridge ID of the currently elected regional root bridge, inside the MSTP region of this bridge. (For the CIST instance only).
Internal Root Cost	The Regional Root Path Cost. For the Regional Root Bridge this is zero. For all other CIST instances in the same MSTP region, it is the sum of the Internal Port Path Costs on the least cost path to the Internal Root Bridge. (For the CIST instance only).
Topology Flag	The current state of the Topology Change Flag of this Bridge instance.
Topology Change Count	he number of times where the topology change flag has been set (during a one-second interval).
Topology Change Last	The time passed since the Topology Flag was last set.
Port	The switch port number of the logical STP port.
Port ID	The port id as used by the STP protocol. This is the priority part and the logical port index of the bridge port.
Role	The current STP port role. The port role can be one of the following values: **`AlternatePort`**, **`BackupPort`**, **`RootPort`**, **`DesignatedPort`**.
State	The current STP port state. The port state can be one of the following values: **`Discarding`**, **`Learning`**, **`Forwarding`**.
Path Cost	The current STP port path cost. This will either be a value computed from the **`Auto`** setting, or any explicitly configured value.
Edge	The current STP port (operational) Edge Flag. An Edge Port is a switch port to which no Bridges are attached. The flag may be automatically computed or explicitly configured. Each Edge Port transits directly to the Forwarding Port State, since there is no possibility of it participating in a loop.
Point-to-Point	The current STP port point-to-point flag. A point-to-point port connects to a non-shared LAN media. The flag may be automatically computed or explicitly configured. The point-to-point properties of a port affect how fast it can transit to STP state.
Uptime	The time since the bridge port was last initialized.

5.15.2 Port Status

This shows the STP CIST port status for physical ports of the switch.
Spanning Tree can be configured at Configuration→Spanning Tree.

STP Port Status

Port	CIST Role	CIST State	Uptime
1	Disabled	Discarding	-
2	Disabled	Discarding	-
3	Disabled	Discarding	-
4	Disabled	Discarding	-
5	Disabled	Discarding	-
6	Disabled	Discarding	-
7	Disabled	Discarding	-
8	RootPort	Forwarding	1d 05:30:20
9	Disabled	Discarding	-
10	Disabled	Discarding	-
11	Disabled	Discarding	-
12	Disabled	Discarding	-

Parameter	Description
Port	The switch port number of the logical STP port.
CIST Role	The current STP port role of the CIST port. The port role can be one of the following values: **AlternatePort**, **BackupPort**, **RootPort**, **DesignatedPort**, **Disabled**.
CIST State	The current STP port role of the CIST port. The port role can be one of the following values: **Discarding**, **Learning**, **Forwarding**.
Uptime	The time since the bridge port was last initialized.

5.15.3 Port Statistics

This shows the STP port statistics counters of bridge ports in the switch.
Spanning Tree can be configured at Configuration→Spanning Tree.

STP Statistics

Port	Transmitted				Received				Discarded	
	MSTP	RSTP	STP	TCN	MSTP	RSTP	STP	TCN	Unknown	Illegal
8	6	0	0	0	0	0	60201	0	0	0

Parameter	Description
Port	The switch port number of the logical STP port.
MSTP	The number of MSTP BPDU's received/transmitted on the port.
RSTP	The number of RSTP BPDU's received/transmitted on the port.
STP	The number of legacy STP Configuration BPDU's received/transmitted on the port.
TCN	The number of (legacy) Topology Change Notification BPDU's received/transmitted on the port.
Discarded Unknown	The number of unknown Spanning Tree BPDU's received (and discarded) on the port.
Discarded Illegal	The number of illegal Spanning Tree BPDU's received (and discarded) on the port.

5.16 MVR

Multiple Registration Protocol (MRP) is a generic Layer 2 framework that supports the dynamic registration and de-registration of attributes, such as VLAN identifiers and multicast MAC addresses, across a bridged LAN, originally defined in IEEE 802.1ak and now incorporated into IEEE 802.1Q. Multiple VLAN Registration Protocol (MVRP) is an MRP application that specifically handles VLAN ID registration, enabling switches to automatically register and withdraw VLAN membership information on trunk interfaces and thus maintain VLAN configuration dynamically across the network. Multicast VLAN Registration (MVR) is a separate Layer 2 multicast feature that allows multicast streams in a dedicated multicast VLAN to be shared with subscriber VLANs so that hosts in different VLANs can receive the same multicast traffic without creating duplicate streams in the core, thereby saving bandwidth

5.16.1 Statistics

This shows MVR Statistics information.
MVR can be configured at Configuration→MVR.

MVR Statistics

VLAN ID	IGMP/MLD Queries Received	IGMP/MLD Queries Transmitted	IGMPv1 Joins Received	IGMPv2/MLDv1 Reports Received	IGMPv3/MLDv2 Reports Received	IGMPv2/MLDv1 Leaves Received
No more entries						

Parameter	Description
VLAN ID	The Multicast VLAN ID.
IGMP/MLD Queries Received	The number of Received Queries for IGMP and MLD, respectively.
IGMP/MLD Queries Transmitted	The number of Transmitted Queries for IGMP and MLD, respectively.
IGMPv1 Joins Received	The number of Received IGMPv1 Join's.
IGMPv2/MLDv1 Report's Received	The number of Received IGMPv2 Join's and MLDv1 Report's, respectively.
IGMPv3/MLDv2 Report's Received	The number of Received IGMPv1 Join's and MLDv2 Report's, respectively.
IGMPv2/MLDv1 Leave's Received	The number of Received IGMPv2 Leave's and MLDv1 Done's, respectively.

5.16.2 IGMP Groups

This shows MVR's IFMP registered multicast entries sorted by VLAN ID, group address, port interface, and source address. Each group always has a so-called ASM (Any-Source Multicast) entry. The ASM entry can be thought of as a catch-remaining entry.
MVR can be configured at Configuration→IPMC Profile and at Configuration→MVR.

MVR IGMP Group Status

VLAN ID	Group Address	Interface	Source Address	Filter Mode	Forwarding	Group Timeout	Source Timeout	In H/W
No Groups								

Parameter	Description
VLAN ID	VLAN ID of the group.
Groups Address	Group ID of the group displayed.
Interface	The port interface for which this entry pertains. If the VLAN ID, Group Address and Interface are identical to the table row above, this is blank.
Source Address	The IP address of the multicast sender. See discussion of ASM entries above for details. Currently, the maximum number of IPv4 source addresses per VLAN ID per group is 8.
Filter Mode	Each <VLAN ID, Group, Interface, Source Address> tuple has a filter mode of either Include or Exclude.
Forwarding	Indicates with a Yes that multicast data is forwarded (allowed) to this port interface or with a No that it is blocked (denied).
Group Timeout	Holds the number of seconds until this group times out. It is only used when Filter Mode is Exclude.
Source Timeout	Holds the number of seconds until this source times out. It is only active when Filter Mode is included.
In H/W	This indicates whether the switch chip has a corresponding entry for this group or <VLAN ID, Group Address, Source Address>. If it says 'No', it is because the switch chip's resources are depleted or

	because the entry is not needed, because the entry is in the so-called "Requested List", because the ASM entry takes care of forwarding.

5.16.3 MLD Groups

This shows MVR's MLD registered multicast entries sorted by VLAN ID, group address, port interface, and source address. Each group always has a so-called ASM (Any-Source Multicast) entry. The ASM entry can be thought of as a catch-remaining entry.
MVR can be configured at Configuration→IPMC Profile and at Configuration→MVR.

MVR MLD Group Status

VLAN ID	Group Address	Interface	Source Address	Filter Mode	Forwarding	Group Timeout	Source Timeout	In H/W
No Groups								

Parameter	Description
VLAN ID	VLAN ID of the group.
Groups Address	Group ID of the group displayed.
Interface	The port interface for which this entry pertains. If the VLAN ID, Group Address and Interface are identical to the table row above, this is blank.
Source Address	The IP address of the multicast sender. See discussion of ASM entries above for details. Currently, the maximum number of IPv6 source addresses per VLAN ID per group is 8.
Filter Mode	Each <VLAN ID, Group, Interface, Source Address> tuple has a filter mode of either Include or Exclude.
Forwarding	Indicates with a Yes that multicast data is forwarded (allowed) to this port interface or with a No that it is blocked (denied).
Group Timeout	Holds the number of seconds until this group times out. It is only used when Filter Mode is Exclude.
Source Timeout	Holds the number of seconds until this source times out. It is only active when Filter Mode is included.
In H/W	This indicates whether the switch chip has a corresponding entry for this group or <VLAN ID, Group Address, Source Address>. If it says 'No', it is because the switch chip's resources are depleted or because the entry is not needed, because the entry is in the so-called "Requested List", because the ASM entry takes care of forwarding.

5.17 IPMC

IPMC is an acronym for **IP M**ulti**C**ast. IPMC supports IPv4 and IPv6 multicasting. IPMCv4 denotes multicast for IPv4. IPMCv6 denotes multicast for IPv6.

5.17.1 IGMP Snooping

5.17.1.1 Status

This shows IGMP Snooping status
IGMP Snooping can be configured at Configuration→IPMC Profile and at Configuration→IPMC→IGMP Snooping.

IGMP Snooping Status

Statistics

VLAN ID	Querier Version	Host Version	Querier Status	Queries Transmitted	Queries Received	V1 Reports Received	V2 Reports Received	V3 Reports Received	V2 Leaves Received

Router Port

Port	Status
1	-
2	-
3	-
4	-
5	-
6	-
7	-
8	-
9	-
10	-
11	-
12	-

Parameter	Description
VLAN ID	VLAN ID of the entry.
Querier Version	Working Querier Version currently.
Host Version	Working Host Version currently.
Querier Status	Shows the Querier status is "ACTIVE" or "IDLE". "DISABLE" denotes the specific interface is administratively disabled.
Queries Transmitted	The number of Transmitted Queries.
Queries Received	The number of Received Queries.
V1 Reports Received	The number of Received V1 Reports.
V2 Reports Received	The number of Received V2 Reports.
V3 Reports Received	The number of Received V3 Reports.
V2 Leaves Received	The number of Received V2 Leaves.
Router Port	Display which ports act as router ports. A router port is a port on the Ethernet switch that leads towards the Layer 3 multicast device or IGMP querier. Static denotes the specific port is configured to be a router port. Dynamic denotes the specific port is learnt to be a router port. Both denote the specific port is configured or learnt to be a router port.
Port	Switch port number.
Status	Indicate whether specific port is a router port or not.

5.17.1.2 Groups

This shows IGMP Group Table. The IGMP Group Table is sorted first by VLAN ID, group address, port interface, and source address. Each group always has a so-called ASM (Any-Source Multicast) entry. The ASM entry can be thought of as a catch-remaining entry.

IGMP Snooping can be configured at Configuration→IPMC Profile and at Configuration→IPMC→IGMP Snooping.

IGMP Snooping Group Status

VLAN ID	Group Address	Interface	Source Address	Filter Mode	Forwarding	Group Timeout	Source Timeout	In H/W
No Groups								

Parameter	Description
VLAN ID	VLAN ID of the group.
Group	Group address of the group displayed.
Port Members	Ports under this group.

5.17.2 MLD Snooping

MLD is an acronym for **M**ulticast **L**istener **D**iscovery for IPv6. MLD is used by IPv6 routers to discover multicast listeners on a directly attached link, much as IGMP is used in IPv4. The protocol is embedded in ICMPv6 instead of using a separate protocol.

5.17.2.1 Status

This shows MLD Snooping status.
MLD Snooping can be configured at Configuration→IPMC Profile and at Configuration→IPMC→MLD Snooping.

MLD Snooping Status

Statistics

VLAN ID	Querier Version	Host Version	Querier Status	Queries Transmitted	Queries Received	V1 Reports Received	V2 Reports Received	V1 Leaves Received

Router Port

Port	Status
1	-
2	-
3	-
4	-
5	-
6	-
7	-
8	-
9	-
10	-
11	-
12	-

Parameter	Description
VLAN ID	VLAN ID of the entry.
Querier Version	Working Querier Version currently.
Host Version	Working Host Version currently.
Querier Status	Shows the Querier status is "ACTIVE" or "IDLE". "DISABLE" denotes the specific interface is administratively disabled.

Queries Transmitted	The number of Transmitted Queries.
Queries Received	The number of Received Queries.
V1 Reports Received	The number of Received V1 Reports.
V2 Reports Received	The number of Received V2 Reports.
V1 Leaves Received	The number of Received V1 Leaves.
Router Port	Display which ports act as router ports. A router port is a port on the Ethernet switch that leads towards the Layer 3 multicast device or MLD querier. Static denotes the specific port is configured to be a router port. Dynamic denotes the specific port is learnt to be a router port. Both denote the specific port is configured or learnt to be a router port.
Port	Switch port number.
Status	Indicate whether specific port is a router port or not.

5.17.2.2 Groups

This shows the MLD registered multicast entries sorted by VLAN ID, group address, port interface, and source address. Each group always has a so-called ASM (Any-Source Multicast) entry. The ASM entry can be thought of as a catch-remaining entry.
MLD Snooping can be configured at Configuration→IPMC Profile and at Configuration→IPMC→MLD Snooping.

MLD Snooping Group Status

VLAN ID	Group Address	Interface	Source Address	Filter Mode	Forwarding	Group Timeout	Source Timeout	In H/W
No Groups								

Parameter	Description
VLAN ID	VLAN ID of the group.
Group Address	Group address of the group displayed.
Interface	The port interface for which this entry pertains. If the VLAN ID, Group Address and Interface are identical to the table row above, this is blank.
Source Address	The IP address of the multicast sender. See discussion of ASM entries above for details. Currently, the maximum number of IPv6 source addresses per VLAN ID per group is 8.
Filter Mode	Each <VLAN ID, Group, Interface, Source Address> tuple has a filter mode of either Include or Exclude.
Forwarding	Indicates with a Yes that multicast data is forwarded (allowed) to this port interface or with a No that it is blocked (denied).
Group Timeout	Holds the number of seconds until this group times out. It is only used when Filter Mode is Exclude.
Source Timeout	Holds the number of seconds until this source times out. It is only active when Filter Mode is include.
In H/W	This indicates whether the switch chip has a corresponding entry for this group or <VLAN ID, Group Address, Source Address>. If it says 'No', it is because the switch chip's resources are depleted or because the entry is not needed, because the entry is in the so-called "Requested List", because the ASM entry takes care of forwarding.

5.18 LLDP

LLDP is an IEEE 802.1ab standard protocol. The Link Layer Discovery Protocol (LLDP) specified in this standard allows stations attached to an IEEE 802 LAN to advertise, to other stations attached to the same IEEE 802 LAN, the major capabilities provided by the system incorporating that station, the management address or addresses of the entity or entities that provide management of those capabilities, and the identification of the stations point of attachment to the IEEE 802 LAN required by those management entity or entities. The information distributed via this protocol is stored by its recipients in a standard Management Information Base (MIB), making it possible for the information to be accessed by a Network Management System (NMS) using a management protocol such as the Simple Network Management Protocol (SNMP).

5.18.1 Neighbors

This shows a status overview for all LLDP neighbors. The displayed table contains a row for each interface on which an LLDP neighbor is detected. The columns hold the following information.
LLDP can be configured at Configuration→LLDP.

LLDP Neighbor Information

LLDP Remote Device Summary						
Local Interface	Chassis ID	Port ID	Port Description	System Name	System Capabilities	Management Address
GigabitEthernet 1/8	00-0B-04-14-7A-74	GigabitEthernet1/0/20	26	9005-AC-C	Bridge(+)	192.168.203.101 (IPv4) - sys-port:20

Parameter	Description
Local Interface	The interface on which the LLDP frame was received.
Chassis ID	The **Chassis ID** is the identification of the neighbor's LLDP frames.
Port ID	The **Port ID** is the identification of the neighbor port.
Port Description	**Port Description** is the port description advertised by the neighbor unit.
System Name	**System Name** is the name advertised by the neighbor unit.
System Capabilities	**System Capabilities** describes the neighbor unit's capabilities. The possible capabilities are: 1. Other 2. Repeater 3. Bridge 4. WLAN Access Point 5. Router 6. Telephone 7. DOCSIS cable device 8. Station only 9. Reserved When a capability is enabled, the capability is followed by (+). If the capability is disabled, the capability is followed by (-).
Management Address	**Management Address** is the neighbor unit's address that is used for higher layer entities to assist discovery by the network management. This could for instance hold the neighbor's IP address.

5.18.2 LLDP-MED Neighbors

This shows a status overview of all LLDP-MED neighbors. LLDP-MED is an extension of IEEE 802.1ab and is Defined by the telecommunication industry association (TIA-1057). The displayed table contains a row for each interface on which an LLDP neighbor is detected. This function applies to VoIP devices which support LLDP-MED. The columns hold the following information:

LLDP-MED can be configured at Configuration→LLDP.

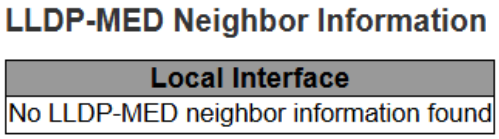

Parameter	Description
Interface	The interface on which the LLDP frame was received.
Device Type	LLDP-MED Devices are comprised of two primary **Device Types:** Network Connectivity Devices and Endpoint Devices. **LLDP-MED Network Connectivity Device Definition** LLDP-MED Network Connectivity Devices, as defined in TIA-1057, provide access to the IEEE 802 based LAN infrastructure for LLDP-MED Endpoint Devices. An LLDP-MED Network Connectivity Device is a LAN access device based on any of the following technologies: 1. LAN Switch/Router 2. IEEE 802.1 Bridge 3. IEEE 802.3 Repeater (included for historical reasons) 4. IEEE 802.11 Wireless Access Point 5. Any device that supports the IEEE 802.1AB and MED extensions defined by TIA-1057 and can relay IEEE 802 frames via any method. **LLDP-MED Endpoint Device Definition** LLDP-MED Endpoint Devices, as defined in TIA-1057, are located at the IEEE 802 LAN network edge, and participate in IP communication service using the LLDP-MED framework. Within the LLDP-MED Endpoint Device category, the LLDP-MED scheme is broken into further Endpoint Device Classes, as defined in the following. Each LLDP-MED Endpoint Device Class is defined to build upon the capabilities defined for the previous Endpoint Device Class. For-example will any LLDP-MED Endpoint Device claiming compliance as a Media Endpoint (Class II) also support all aspects of TIA-1057 applicable to Generic Endpoints (Class I), and any LLDP-MED Endpoint Device claiming compliance as a Communication Device (Class III) will also support all aspects of TIA-1057 applicable to both Media Endpoints (Class II) and Generic Endpoints (Class I). **LLDP-MED Generic Endpoint (Class I)** The LLDP-MED Generic Endpoint (Class I) definition is applicable to all endpoint products that require the base LLDP discovery services defined in TIA-1057, however do not support IP media or act as an end-user communication appliance. Such devices may include (but are not limited to) IP Communication Controllers, other communication related servers, or any device requiring basic services as defined in TIA-1057. Discovery services defined in this class include LAN configuration, device location, network policy, power management, and inventory management. **LLDP-MED Media Endpoint (Class II)** The LLDP-MED Media Endpoint (Class II) definition is applicable to all endpoint products that have IP media capabilities however may or may not be associated with a particular end user. Capabilities include all of the capabilities defined for the previous Generic Endpoint Class (Class I), and are extended to include aspects related to media streaming. Example product categories expected to adhere to this class include (but are not limited to) Voice / Media Gateways, Conference Bridges, Media Servers, and similar. Discovery services defined in this class include media-type-specific network layer policy discovery. **LLDP-MED Communication Endpoint (Class III)** The LLDP-MED Communication Endpoint (Class III) definition is applicable to all endpoint products that act as end user communication appliances supporting IP media. Capabilities include all of the capabilities defined for the previous Generic Endpoint (Class I) and Media Endpoint (Class II) classes, and are extended to include aspects related to end user devices. Example product categories expected to adhere to this class include (but are not limited to) end user communication appliances, such as IP Phones, PC-based softphones, or other communication appliances that directly support the end user. Discovery services defined in this class include provision of location identifier (including ECS / E911 information), embedded L2 switch support, inventory management.
LLDP-MED Capabilities	**LLDP-MED Capabilities** describes the neighbor unit's LLDP-MED capabilities. The possible capabilities are:

	1. LLDP-MED capabilities 2. Network Policy 3. Location Identification 4. Extended Power via MDI - PSE 5. Extended Power via MDI - PD 6. Inventory 7. Reserved
Application Type	**Application Type** indicating the primary function of the application(s) defined for this network policy, advertised by an Endpoint or Network Connectivity Device. The possible application types are shown below. 1. Voice - for use by dedicated IP Telephony handsets and other similar appliances supporting interactive voice services. These devices are typically deployed on a separate VLAN for ease of deployment and enhanced security by isolation from data applications. 2. Voice Signaling - for use in network topologies that require a different policy for the voice signaling than for the voice media. 3. Guest Voice - to support a separate limited feature-set voice service for guest users and visitors with their own IP Telephony handsets and other similar appliances supporting interactive voice services. 4. Guest Voice Signaling - for use in network topologies that require a different policy for the guest voice signaling than for the guest voice media. 5. Softphone Voice - for use by softphone applications on typical data centric devices, such as PCs or laptops. 6. Video Conferencing - for use by dedicated Video Conferencing equipment and other similar appliances supporting real-time interactive video/audio services. 7. Streaming Video - for use by broadcast or multicast based video content distribution and other similar applications supporting streaming video services that require specific network policy treatment. Video applications relying on TCP with buffering would not be an intended use of this application type. 8. Video Signaling - for use in network topologies that require a separate policy for the video signaling than for the video media.
Policy	**Policy** indicates that an Endpoint Device wants to explicitly advertise that the policy is required by the device. Can be either Defined or Unknown Unknown: The network policy for the specified application type is currently unknown. Defined: The network policy is defined (known).
TAG	**TAG** is indicative of whether the specified application type is using a tagged or an untagged VLAN. Can be Tagged or Untagged. Untagged: The device is using an untagged frame format and as such does not include a tag header as defined by IEEE 802.1Q-2003. Tagged: The device is using the IEEE 802.1Q tagged frame format.
VLAN ID	**VLAN ID** is the VLAN identifier (VID) for the interface as defined in IEEE 802.1Q-2003. A value of 1 through 4094 is used to define a valid VLAN ID. A value of 0 (Priority Tagged) is used if the device is using priority tagged frames as defined by IEEE 802.1Q-2003, meaning that only the IEEE 802.1D priority level is significant and the default PVID of the ingress interface is used instead.
Priority	**Priority** is the Layer 2 priority to be used for the specified application type. One of the eight priority levels (0 through 7).
DSCP	**DSCP** is the DSCP value to be used to provide Diffserv node behavior for the specified application type as defined in IETF RFC 2474. Contain one of 64 code point values (0 through 63).
Auto-negotiation	**Auto-negotiation** identifies if MAC/PHY auto-negotiation is supported by the link partner.
Auto-negotiation status	**Auto-negotiation status** identifies if auto-negotiation is currently enabled at the link partner. If **Auto-negotiation** is supported and **Auto-negotiation status** is disabled, the 802.3 PMD operating mode will be determined the operational MAU type field value rather than by auto-negotiation.
Auto-negotiation Capabilities	**Auto-negotiation Capabilities** shows the link partners MAC/PHY capabilities.

5.18.3 PoE

Power budget note: the monitor page reports PoE delivery status and consumption. For 9561-8GP4XS-TSN, total PD delivery is limited to 360 W even though the PSU peak capability is 400 W.

This shows a status overview for all LLDP PoE neighbors. The displayed table contains a row for each interface on which an LLDP PoE neighbor is detected. The columns hold the following information.
LLDP PoE can be configured at Configuration→LLDP.

LLDP Neighbor Power Over Ethernet Information

Local Interface	Power Type	Power Source	Power Priority	Maximum Power
GigabitEthernet 1/8	Reserved	PSE	Low	0 [W]

Parameter	Description
Local Interface	The interface on which the LLDP frame was received.
Power Type	The **Power Type** represents whether the device is a Power Sourcing Entity (PSE) or Power Device (PD). If the **Power Type** is unknown, it is represented as "Reserved".
Power Source	The **Power Type** represents the power source being utilized by a PSE or PD device. If the device is a PSE device it can either run on its Primary Power Source or its Backup Power Source. If it is unknown whether the PSE device is using its Primary Power Source or its Backup Power Source it is indicated as "Unknown" If the device is a PD device it can either run on its local power supply or it can use the PSE as power source. It can also use both its local power supply and the PSE. If it is unknown what power supply the PD device is using it is indicated as "Unknown"
Power Priority	Power **Power Priority** represents the priority of the PD device, or the power priority associated with the PSE type device's interface that is sourcing the power. There are three levels of power priority. The three levels are: Critical, High and Low. If the power priority is unknown, it is indicated as "Unknown"
Maximum Power	The **Maximum Power** Value contains a numerical value that indicates the maximum power in watts required by a PD device from a PSE device, or the minimum power a PSE device is capable of sourcing over a maximum length cable based on its current configuration. The maximum allowed value is 102.3 W. If the device indicates value higher than 102.3 W, it is represented as "reserved"

Figure: LLDP PoE neighbor monitoring page.

LLDP Neighbor Power Over Ethernet Information

Auto-refresh Refresh

Local Interface	Power Type	Power Source	Power Priority	Maximum Power
No PoE neighbor information found				

5.18.4 EEE

This shows an overview of EEE information exchanged by LLDP. EEE is an abbreviation for Energy Efficient Ethernet defined in IEEE 802.3az. By using EEE power savings can be achieved at the expense of traffic latency. This latency occurs due to the EEE circuits turning off to save power, need time to boot up before sending traffic over the link. This time is called “wakeup time”. To achieve minimal latency, devices can use LLDP to exchange information about their respective tx and rx “wakeup time”, as a way to agree upon the minimum wakeup time they need.
LLDP can be configured at Configuration→LLDP.

LLDP Neighbors EEE information:
If the interface does not support EEE, then it displays as **`EEE not supported for this interface`**.
If EEE is not enabled on this particular interface, then it displays as **`EEE not enabled for this interface`**.
If the link partner doesn't support EEE, then it displays as **`Link partner is not EEE capable`**.

LLDP Neighbors EEE Information

Local Interface	Tx Tw	Rx Tw	Fallback Receive Tw	Echo Tx Tw	Echo Rx Tw	Resolved Tx Tw	Resolved Rx Tw	EEE in Sync
GigabitEthernet 1/8	EEE not supported for this interface							

Parameter	Description
Local Interface	The interface at which LLDP frames are received or transmitted.
Tx Tw	The link partner's maximum time that transmits path can hold-off sending data after deassertion of LPI.
Rx Tw	The link partner's time that receiver would like the transmitter to hold-off to allow time for the receiver to wake from sleep.
Fallback Receive Tw	The link partner's fallback receives Tw. A receiving link partner may inform the transmitter of an alternate desired Tw_sys_tx. Since a receiving link partner is likely to have discrete levels for savings, this provides the transmitter with additional information that it may use for a more efficient allocation. Systems that do not implement this option default the value to be the same as that of the Receive Tw_sys_tx.

Echo Tx Tw	The link partner's Echo Tx Tw value. The respective echo values shall be defined as the local link partners reflection (echo) of the remote link partners respective values. When a local link partner receives its echoed values from the remote link partner it can determine whether or not the remote link partner has received, registered and processed its most recent values. For example, if the local link partner receives echoed parameters that do not match the values in its local MIB, then the local link partner infers that the remote link partners request was based on stale information.
Echo Rx Tw	The link partner's Echo Rx Tw value.
Resolved Tx Tw	The resolved Tx Tw for this link. Note: NOT the link partner The resolved value that is the actual "tx wakeup time " used for this link (based on EEE information exchanged via LLDP).
Resolved Rx Tw	The resolved Rx Tw for this link. Note: NOT the link partner The resolved value that is the actual "tx wakeup time " used for this link (based on EEE information exchanged via LLDP).
EEE in Sync	Shows whether the switch and the link partner have agreed on wake times. Red - Switch and link partner have not agreed on wakeup times. Green - Switch and link partner have agreed on wakeup times.

5.18.5 Port Statistics

This shows an overview of all LLDP traffic. Two types of counters are shown. **Global counters** are counters that refer to the whole switch, while **local counters** refer to per interface counters for the currently selected switch. LLDP can be configured at Configuration→LLDP.

LLDP Global Counters

Global Counters	
Clear global counters	☑
Neighbor entries were last changed	2025-12-20T03:25:28+00:00 (1 secs. ago)
Total Neighbors Entries Added	2
Total Neighbors Entries Deleted	1
Total Neighbors Entries Dropped	0
Total Neighbors Entries Aged Out	1

LLDP Statistics Local Counters

Local Interface	Tx Frames	Rx Frames	Rx Errors	Frames Discarded	TLVs Discarded	TLVs Unrecognized	Org. Discarded	Age-Outs	Clear
*	*	*	*	*	*	*	*	*	☑
GigabitEthernet 1/1	0	0	0	0	0	0	0	0	☑
GigabitEthernet 1/2	0	0	0	0	0	0	0	0	☑
GigabitEthernet 1/3	0	0	0	0	0	0	0	0	☑
GigabitEthernet 1/4	0	0	0	0	0	0	0	0	☑
GigabitEthernet 1/5	0	0	0	0	0	0	0	0	☑
GigabitEthernet 1/6	0	0	0	0	0	0	0	0	☑
GigabitEthernet 1/7	0	0	0	0	0	0	0	0	☑
GigabitEthernet 1/8	33700	67364	0	0	0	471564	336820	1	☑
10GigabitEthernet 1/1	0	0	0	0	0	0	0	0	☑
10GigabitEthernet 1/2	0	0	0	0	0	0	0	0	☑
10GigabitEthernet 1/3	0	0	0	0	0	0	0	0	☑
10GigabitEthernet 1/4	0	0	0	0	0	0	0	0	☑

Parameter	Description
Clear global counters	If checked the global counters are cleared when ”Clear” is pressed.
Neighbor entries were last changed	Shows the time when the last entry was last deleted or added. It also shows the time elapsed since the last change was detected.
Total Neighbors Entries Added	Shows the number of new entries added since switch reboot.
Total Neighbors Entries Deleted	Shows the number of new entries deleted since switch reboot.

Total Neighbors Entries Dropped	Shows the number of LLDP frames dropped due to the entry table being full.
Total Neighbors Entries Aged Out	Shows the number of entries deleted due to Time-To-Live expiring.
Local Interface	The interface on which LLDP frames are received or transmitted.
Tx Frames	The number of LLDP frames transmitted on the interface.
Rx Frames	The number of LLDP frames received on the interface.
Rx Errors	The number of received LLDP frames containing some kind of error.
Frames Discarded	If a LLDP frame is received on an interface, and the switch's internal table has run full, the LLDP frame is counted and discarded. This situation is known as "Too Many Neighbors" in the LLDP standard. LLDP frames require a new entry in the table when the Chassis ID or Remote Port ID is not already contained within the table. Entries are removed from the table when a given interface's link is down, an LLDP shutdown frame is received, or when the entry ages out.
TLVs Discarded	Each LLDP frame can contain multiple pieces of information, known as TLVs (TLV is short for "Type Length Value"). If a TLV is malformed, it is counted and discarded.
TLVs Unrecognized	The number of well-formed TLVs, but with an unknown type value.
Org. Discarded	If LLDP frame is received with an organizationally TLV, but the TLV is not supported the TLV is discarded and counted.
Age-Outs	Each LLDP frame contains information about how long time the LLDP information is valid (age-out time). If no new LLDP frame is received within the age out time, the LLDP information is removed, and the **Age-Out** counter is incremented.
Clear	If checked the counters are cleared when ”Clear” is pressed.

5.19 PTP

Precision Time Protocol (PTP) is a network protocol for synchronizing the clocks of computer systems.

5.19.1 PTP

This page allows the user to inspect the current PTP clock settings.
PTP can be configured at Configuration→PTP.

PTP External Clock Mode

One_PPS_Mode	Output
External Enable	False
Adjust Method	Auto
Clock Frequency	1
One PPS Domain	0

PTP Clock Configuration

			Port List										
Inst	**ClkDom**	**Device Type**	1	2	3	4	5	6	7	8	9	10	11	12
No Clock Instances Present														

Parameter	Description
One_PPS_Mode	Shows the current One_pps_mode configured. 1. Output: Enable the 1 pps clock output

	2. Disable: Disable the 1 pps clock output
External Enable	Shows the current External clock output configuration. 1. True: Enable the external clock output 2. False: Disable the external clock output
Adjust Method	Shows the current Frequency adjustment configuration. 1. LTC: Use Local Time Counter (LTC) frequency control 2. Auto: AUTO Select clock control, based on PTP profile and available HW resources.
Clock Frequency	Shows the current clock frequency used by the External Clock. The possible range of values are 1 - 25000000 (1 - 25MHz)
One PPS Domain	Shows the current clock domain used for generating 1pps clock output.
Inst	Indicates the Instance of a particular Clock Instance [0..3]. Click on the Clock Instance number to monitor the Clock details.
ClkDom	Indicates the Clock domain used by the Instance of a particular Clock Instance [0..3].
Device Type	Indicates the Type of the Clock Instance. There are five Device Types. 1. Ord-Bound - Clock's Device Type is Ordinary-Boundary Clock. 2. P2p Transp - Clock's Device Type is Peer to Peer Transparent Clock. 3. E2e Transp - Clock's Device Type is End to End Transparent Clock. 4. Master Only - Clock's Device Type is Master Only. 5. Slave Only - Clock's Device Type is Slave Only. 6. AED GM - clock's Device Type is Grandmaster only (AED profile specific).
Port List	Shows the ports configured for that Clock Instance.

5.19.2 PTP Statistics

This page allows the user to inspect the current PTP counters. PTP can be configured at Configuration→PTP.

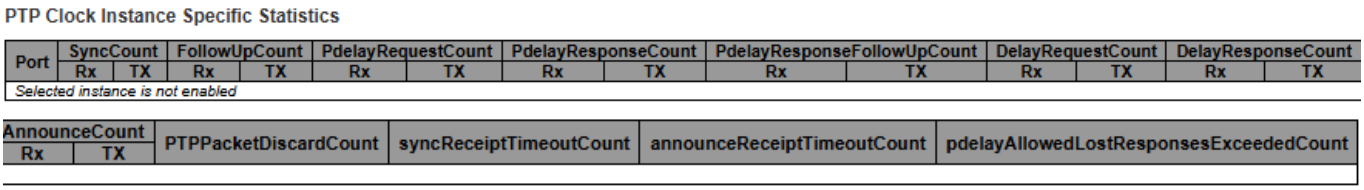

Parameter	Description
Instance	PTP Instance (Range 0-3), CMLDS (Common Mean Link Delay Service).
SyncCount	A counter that increments every time when synchronization information is received.
FollowUpCount	A counter that increments every time when a Follow Up message is received.
PdelayRequestCount	A counter that increments every time when a Pdelay_Req message is received.
PdelayResponseCount	A counter that increments every time when a Pdelay_Resp message is received.
PdelayResponseFollowUpCount	A counter that increments every time when a Pdelay_Resp_Follow_Up message is received.
DelayRequestCount	A counter that increments every time when a Delay_Req message is received.
DelayResponseCount	A counter that increments every time when a Delay_Resp message is received.
AnnounceCount	A counter that increments every time when an Announce message is received.
PTPPacketDiscardCount	A counter that increments every time when a PTP message is discarded.

syncReceiptTimeoutCount	A counter that increments every time when sync receipt timeout occurs.
announceReceiptTimeoutCount	A counter that increments every time when announce receipt timeout occurs.
pdelayAllowedLostResponsesExceededCount	A counter that increments everytime the value of the variable lostResponses exceeds the value of the variable allowedLostResponses.
SyncCount	A counter that increments every time synchronization information is transmitted
FollowUpCount	A counter that increments every time a Follow_Up message is transmitted.
PdelayRequestCount	A counter that increments every time a Pdelay_Req message is transmitted.
PdelayResponseCount	A counter that increments every time a Pdelay_Resp message is transmitted.
PdelayResponseFollowUpCount	A counter that increments every time a Pdelay_Resp_Follow_Up message is transmitted.
DelayRequestCount	A counter that increments every time a Delay_Req message is transmitted.
DelayResponseCount	A counter that increments every time a Delay_Resp message is transmitted.
AnnounceCount	A counter that increments every time an Announce message is transmitted.

5.20 PoE

This page allows the user to inspect the current status for all PoE ports. PoE is an acronym for **P**ower **O**ver **E**thernet. Power Over Ethernet is used to transmit electrical power, to remote devices over standard Ethernet cable. It could for example be used for powering IP telephones, wireless LAN access points and other equipment, where it would be difficult or expensive to connect the equipment to main power supply.
PoE can be configured at Configuration→PoE.

Power Over Ethernet Status

Port	PD Class	Power Requested	Power Allocated	Power Used	Current Used	Priority	Port Status
1	-	0 [W]	0 [W]	0 [W]	0 [mA]	Low	---
2	-	0 [W]	0 [W]	0 [W]	0 [mA]	Low	---
3	-	0 [W]	0 [W]	0 [W]	0 [mA]	Low	---
4	-	0 [W]	0 [W]	0 [W]	0 [mA]	Low	---
5	-	0 [W]	0 [W]	0 [W]	0 [mA]	Low	---
6	-	0 [W]	0 [W]	0 [W]	0 [mA]	Low	---
7	-	0 [W]	0 [W]	0 [W]	0 [mA]	Low	---
8	-	0 [W]	0 [W]	0 [W]	0 [mA]	Low	---

Parameter	Description
Port	This is the logical port number for this row.
PD Class	Shows the detected or reported PD power class/type for the connected powered device. Current 9561 PoE port type support is Type3-15W, Type3-30W, Type3-60W, and Type4-90W; unsupported class-only entries such as Class 0, Class 1, Class 2, Class 5, and Class 7 are not listed as configurable PoE port types.
Power Requested	The **Power Requested** shows the requested amount of power the PD wants to be reserved.
Power Allocated	The **Power Allocated** shows the amount of power the switch has allocated for the PD.
Power Used	The **Power Used** shows how much power the PD currently is using.
Current Used	The **Power Used** shows how much current the PD currently is using.

Priority	The **Priority** shows the port's priority configured by the user.
Port Status	The **Port Status** shows the port's status. The status can be one of the following values: **On** – A PD is detected for the port. **---** - No PD detected for the port. **Not Supported** – PoE not supported for the port. **Budget Exceeded** - The total requested or used power by the PDs exceeds the maximum power the Power Supply can deliver, and port(s) with the lowest priority is/are powered down. **Off** – PD is off. **Disabled** - User has disabled PoE for the port. **Shutdown** - The port is shut down. **Overload** – The PD has requested or used more power than the port can deliver, and is powered down. **Unknown** – PD detected, but is not working correctly.

5.21 MAC Table

This shows the MAC Table information. The MAC Table contains up to 32768 entries, and is sorted first by VLAN ID, then by MAC address. Switching of frames is based upon the DMAC address contained in the frame. The switch builds up a table that maps MAC addresses to switch ports for knowing which ports the frames should go to (based upon the DMAC address in the frame). This table contains both static and dynamic entries. The static entries are configured by the network administrator if the administrator wants to do a fixed mapping between the DMAC address and switch ports.

The frames also contain a MAC address (SMAC address), which shows the MAC address of the equipment sending the frame. The SMAC address is used by the switch to automatically update the MAC Table with these dynamic MAC addresses. Dynamic entries are removed from the MAC table if no frame with the corresponding SMAC address have been seen after a configurable age time.

MAC Table can be configured at Configuration→MAC Table.

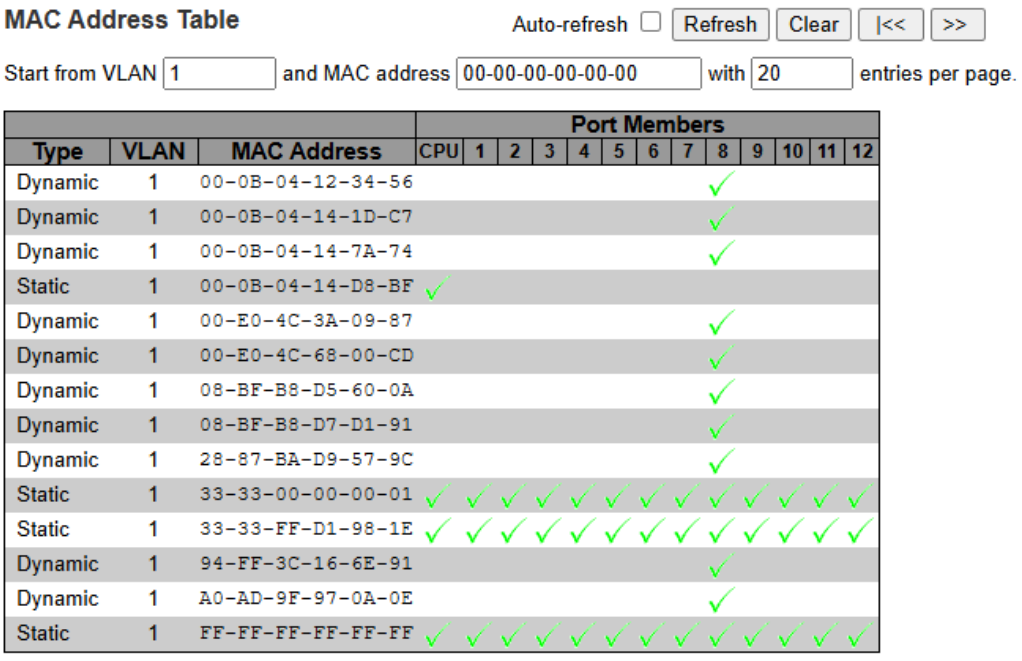

Parameter	Description
Type	Indicates whether the entry is a static or a dynamic entry.
MAC address	The MAC address of the entry.

VLAN	The VLAN ID of the entry.
Port Members	The ports that are members of the entry.
PSFP ID	Per-Stream Filtering and Policing ID for MAC address entry.

5.22 VLANs

Virtual LAN. A method to restrict communication between switch ports. At layer 2, the network is portioned into multiple, distinct, mutually isolated broadcast domains.

5.22.1 Membership

This shows an overview of membership status of VLAN users.
VLAN Membership can be configured at Configuration→VLANs.

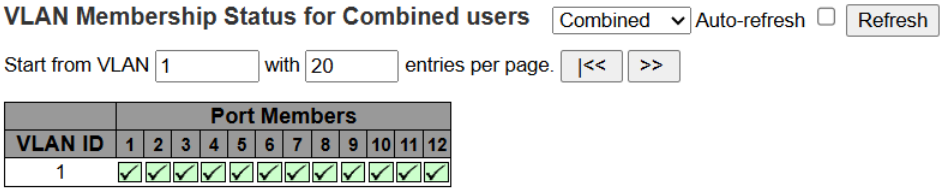

Parameter	Description
VLAN User	Select VLAN Users from the drop-down list. Various internal software modules may use VLAN services to configure VLAN memberships on the fly. The drop-down list on the right allows for selecting between showing VLAN memberships as configured by an administrator (Admin) or as configured by one of these internal software modules. The "Combined" entry will show a combination of the administrator and internal software modules configuration, and basically reflects what is actually configured in hardware.
VLAN ID	VLAN ID for which the Port members are displayed.
Port Members	A row of check boxes for each port is displayed for each VLAN ID. If a port is included in a VLAN, the following image will be displayed: ✓. If a port is in the forbidden port list, the following image will be displayed: ✗. If a port is in the forbidden port list and at the same time attempted included in the VLAN, the following image will be displayed: ✗. The port will not be a member of the VLAN in this case.

5.22.2 Port

This shows VLAN Port Status.
VLAN Port can be configured at Configuration→VLANs.

VOLKTEK

VLAN Port Status for Combined users

Combined | Auto-refresh | Refresh

Port	Port Type	Ingress Filtering	Frame Type	Port VLAN ID	Tx Tag	Untagged VLAN ID	Conflicts
1	C-Port	☑	All	1	Untag All		No
2	C-Port	☑	All	1	Untag All		No
3	C-Port	☑	All	1	Untag All		No
4	C-Port	☑	All	1	Untag All		No
5	C-Port	☑	All	1	Untag All		No
6	C-Port	☑	All	1	Untag All		No
7	C-Port	☑	All	1	Untag All		No
8	C-Port	☑	All	1	Untag All		No
9	C-Port	☑	All	1	Untag All		No
10	C-Port	☑	All	1	Untag All		No
11	C-Port	☑	All	1	Untag All		No
12	C-Port	☑	All	1	Untag All		No

Parameter	Description
VLAN User	Various internal software modules may use VLAN services to configure VLAN port configuration on the fly. The drop-down list on the right allows for selecting between showing VLAN memberships as configured by an administrator (Admin) or as configured by one of these internal software modules. The "Combined" entry will show a combination of the administrator and internal software modules configuration, and basically reflects what is actually configured in hardware. If a given software modules haven't overridden any of the port settings, the text "No data exists for the selected user" is shown in the table.
Port	The logical port for the settings contained in the same row.
Port Type	Shows the port type (Unaware, C-Port, S-Port, S-Custom-Port.) that a given user wants to configure on the port. The field is empty if not overridden by the selected user.
Ingress Filtering	Shows whether a given user wants ingress filtering enabled or not. The field is empty if not overridden by the selected user.
Frame Type	Shows the acceptable frame types (All, Tagged, Untagged) that a given user wants to configure on the port. The field is empty if not overridden by the selected user.
Port VLAN ID	Shows the Port VLAN ID (PVID) that a given user wants the port to have. The field is empty if not overridden by the selected user.
Tx Tag	Shows the Tx Tag requirements (Tag All, Tag PVID, Tag UVID, Untag All, Untag PVID, Untag UVID) that a given user has on a port. The field is empty if not overridden by the selected user.
Untagged VLAN ID	If Tx Tag is overridden by the selected user and is set to Tag or Untag UVID, then this field will show the VLAN ID the user wants to tag or untag on egress. The field is empty if not overridden by the selected user.
Conflicts	Two users may have conflicting requirements to a port's configuration. For instance, one user may require all frames to be tagged on egress while another requires all frames to be untagged on egress. Since both users cannot win, this gives rise to a conflict, which is solved in a prioritized way. The Administrator has the least priority. Other software modules are prioritized according to their position in the drop-down list: The higher in the list, the higher priority. If conflicts exist, it will be displayed as "Yes" for the "Combined" user and the offending software module. The "Combined" user reflects what is actually configured in hardware.

5.23 MVRP

This shows the statistics for the Multiple Vlan Registration Protocol (MVRP) protocol for all switch ports. MVRP is a protocol that defines the dynamic registration and de-registration of VLAN identifiers across a Bridged Local Area Network. It is using the MRP framework to define its operation and therefore it is also called a MRP Application. The standard was originally defined by IEEE 802.1ak, and its latest incorporation is in IEEE 802.1Q-2014.

MVRP can be configured at Configuration→MRP.

MVRP Statistics

Port	Failed Registrations	Last PDU Origin
1	0	00-00-00-00-00-00
2	0	00-00-00-00-00-00
3	0	00-00-00-00-00-00
4	0	00-00-00-00-00-00
5	0	00-00-00-00-00-00
6	0	00-00-00-00-00-00
7	0	00-00-00-00-00-00
8	0	00-00-00-00-00-00
9	0	00-00-00-00-00-00
10	0	00-00-00-00-00-00
11	0	00-00-00-00-00-00
12	0	00-00-00-00-00-00

Parameter	Description
Port	The logical port for the statistics contained in the same row.
Failed Registrations	The number of failed VLAN registrations on this switch port. Each port implementing the MVRP protocol maintains a count of the number of times it has received a VLAN registration request but has failed to register the VLAN due to lack of space in the Filtering Database.
Port Members	The MAC address of the most recent MVRP PDU received on this switch port. MAC is 00-00-00-00-00-00 if the protocol is not enabled on that switch port, or if the port has not received any MVRP PDUs yet.

5.24 sFlow

This shows the receiver and per-port sFlow statistics. sFlow is an industry standard technology for monitoring switched networks through random sampling of packets on switch ports and time-based sampling of port counters. The sampled packets and counters (referred to as flow samples and counter samples, respectively) are sent as sFlow UDP datagrams to a central network traffic monitoring server. This central server is called an sFlow receiver or sFlow collector. Additional information can be found at http://sflow.org.
sFlow can be configured at Configuration→sFlow.

sFlow Statistics Auto-refresh ☐ Refresh Clear Receiver Clear Ports

Receiver Statistics

Owner	<none>
IP Address/Hostname	0.0.0.0
Timeout	0
Tx Successes	0
Tx Errors	0
Flow Samples	0
Counter Samples	0

Port Statistics

Port	Flow Samples	Counter Samples
1	0	0
2	0	0
3	0	0
4	0	0
5	0	0
6	0	0
7	0	0
8	0	0
9	0	0
10	0	0
11	0	0
12	0	0

Parameter	Description
Owner	This field shows the current owner of the sFlow configuration. It assumes one of three values as follows: • If sFlow is currently unconfigured/unclaimed, Owner contains **<none>**.

	• If sFlow is currently configured through Web or CLI, Owner contains **<Configured through local management>.** • If sFlow is currently configured through SNMP, Owner contains a string identifying the sFlow receiver.
IP Address/Hostname	The IP address or hostname of the sFlow receiver.
Timeout	The number of seconds remaining before sampling stops and the current sFlow owner is released.
Tx Successes	The number of UDP datagrams successfully sent to the sFlow receiver.
Tx Errors	The number of UDP datagrams that has failed transmission. The most common source of errors is invalid sFlow receiver IP/hostname configuration. To diagnose, paste the receiver's IP address/hostname into the Ping Web page (Diagnostics → Ping/Ping6).
Flow Samples	The total number of flow samples sent to the sFlow receiver.
Counter Samples	The total number of counter samples sent to the sFlow receiver.
Port	The port number for which the following statistics applies.
Flow Samples	The number of flow samples sent to the sFlow receiver originating from this port.
Counter Samples	The total number of counter samples sent to the sFlow receiver originating from this port.
Clear Receiver	Clears the sFlow receiver counters.
Clear Ports	Clears the per-port counters.

5.25 DDMI

Digital Diagnostics Monitoring Interface (DDMI) provides an enhanced digital diagnostic monitoring interface for optical transceivers which allows real-time access to device operating parameters.

5.25.1 Overview

This shows the DDMI overview information.
DDMI can be configured at Configuration→DDMI.

DDMI Overview

Port	Vendor	Part Number	Serial Number	Revision	Data Code	Transceiver
9	-	-	-	-	-	-
10	-	-	-	-	-	-
11	-	-	-	-	-	-
12	-	-	-	-	-	-

Parameter	Description
Port	DDMI port.
Vendor	Indicates Vendor name SFP vendor name.
Part Number	Indicates Vendor PN Part number provided by SFP vendor.
Serial Number	Indicates Vendor SN Serial number provided by vendor.

Revision	Indicates Vendor rev Revision level for part number provided by vendor.
Data Code	Indicates Date code Vendor's manufacturing date code.
Transceiver	Indicates Transceiver compatibility.

5.25.2 Detailed

This shows the DDMI detail information.
DDMI can be configured at Configuration→DDMI.

Transceiver Information

Port 9 Auto-refresh Refresh

Vendor	-
Part Number	-
Serial Number	-
Revision	-
Data Code	-
Transceiver	-

DDMI Information

Type	Current	Alarm/Warning	Low Warning Threshold	High Warning Threshold	Low Alarm Threshold	High Alarm Threshold
Temperature [C]	-	-	-	-	-	-
Voltage [V]	-	-	-	-	-	-
Tx Bias [mA]	-	-	-	-	-	-
Tx Power [mW]	-	-	-	-	-	-
Rx Power [mW]	-	-	-	-	-	-

Parameter	Description
Vendor	Indicates SFP vendor name.
Part Number	Indicates part number provided by SFP vendor.
Serial Number	Indicates serial number provided by SFP vendor.
Revision	Indicates revision level for part number provided by SFP vendor.
Data Code	Indicates vendor's manufacturing date code.
Transceiver	Indicates SFP transceiver compatibility.
Current	The current value of temperature, voltage, Tx bias, Tx power, and Rx power.
Alarm/Warning	Indicates whether there is an alarm or warning.
Low Warning Threshold	The low warning threshold value of temperature, voltage, Tx bias, Tx power, and Rx power.
High Warning Threshold	The high warning threshold value of temperature, voltage, Tx bias, Tx power, and Rx power.
Low Alarm Threshold	The low alarm threshold value of temperature, voltage, Tx bias, Tx power, and Rx power.
High Alarm Threshold	The high alarm threshold value of temperature, voltage, Tx bias, Tx power, and Rx power.

5.26 UDLD

This shows the UDLD status. UDLD is an acronym for **U**ni **D**irectional **L**ink **D**etection. UDLD protocol monitors the physical configuration of the links between devices and ports that support UDLD. It detects the existence of unidirectional links. Its functionality is to provide mechanisms useful for detecting one-way connections before they create a loop or other protocol malfunction. RFC 5171 specifies a way at data link layer to detect Uni directional link.

UDLD can be configured at Configuration→UDLD.

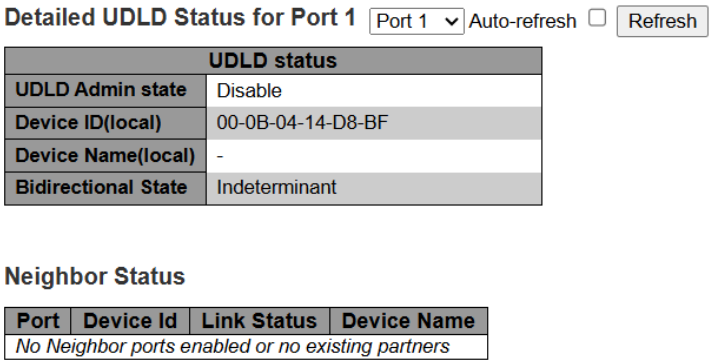

Parameter	Description
UDLD Admin State	The current port state of the logical port, Enabled if any of state (Normal, Aggressive) is Enabled.
Device ID (local)	The ID of Device.
Device Name (local)	Name of the Device.
Bidirectional State	The current state of the port.
Port	The current port of neighbor device.
Device ID	The current ID of neighbor device.
Link Status	The current link status of neighbor port.
Device Name	Name of the Neighbor Device.

5.27 TSN

This section contains status information about TSN.

To configure TSN streams, refer to Configuration→TSN.

5.27.1 Frame Preemption

This shows an overview of TSN Egress Port Frame Preemption Status aal switch ports.

TSN Egress Port Frame Preemption can be configured at Configuration→TSN→Frame Preemption.

TSN Egress Port Frame Preemption Status

Auto-refresh ☐ Refresh

Port	Hold Advance	Release Advance	Preemption Active	Hold Request	Status Verify	LocPreemptsupport	LocPreemptEnabled	LocPreemptActive	LocAddFragSize
1	0	0	✗	✗	indeterminate	✓	✗	✓	0
2	0	0	✗	✗	indeterminate	✓	✗	✓	0
3	0	0	✗	✗	indeterminate	✓	✗	✓	0
4	0	0	✗	✗	indeterminate	✓	✗	✓	0
5	0	0	✗	✗	indeterminate	✓	✗	✓	0
6	0	0	✗	✗	indeterminate	✓	✗	✓	0
7	0	0	✗	✗	indeterminate	✓	✗	✓	0
8	0	0	✗	✗	disabled	✓	✗	✓	0
9	0	0	✗	✗	indeterminate	✓	✗	✓	0
10	0	0	✗	✗	indeterminate	✓	✗	✓	0
11	0	0	✗	✗	indeterminate	✓	✗	✓	0
12	0	0	✗	✗	indeterminate	✓	✗	✓	0

Parameter	Description
Port	The logical port for the settings contained in the same row.
Hold Advance	The value of the holdAdvance parameter for the Port in nanoseconds. There is no default value; the holdAdvance is a property of the underlying MAC.
Release Advance	The value of the releaseAdvance parameter for the Port in nanoseconds. There is no default value; the releaseAdvance is a property of the underlying MAC.
Preemption Active	The value is active (TRUE) when preemption is operationally active for the Port, and idle (FALSE) otherwise.
Hold Request	The value is hold (TRUE) when the sequence of gate operations for the Port has executed a Set-And-Hold-MAC operation, and release (FALSE) when the sequence of gate operations has executed a Set-And-Release-MAC operation. The value of this object is release (FALSE) on system initialization.
Status Verify	The status of the MAC Merge sublayer verification for the given device. The possible values are: • initial: Verification process is initializing • idle: The verification process is ongoing • send: The verification process is ongoing • wait: The verification process is ongoing • succeeded: The verification process is completed and frame preemption is supported • failed: The verification process is completed and frame preemption is not supported • disabled: Frame preemption is disabled • indeterminate: The verification process is completed but it cannot be determined whether link partner supports frame preemption
LocPreemptsupport	The value is TRUE when preemption is supported on the port, and FALSE otherwise.
LocPreemptEnabled	The value is TRUE when preemption is enabled on the port, and FALSE otherwise.
LocPreemptActive	The value is TRUE when preemption is operationally active on the port, and FALSE otherwise.
LocAddFragSize	The value of the 802.3br LocAddFragSize parameter for the port. The minimum size of non-final fragments supported by the receiver on the local port. This value is expressed in units of 64 octets of additional fragment length. The minimum non-final fragment size is: (LocAddFragSize + 1) * 64 octets.

5.27.2 TAS

This shows user how to check the current TAS configuration. TAS is an acronym for **T**ime **A**ware **S**haper. 802.1Qbv: This amendment specifies time-aware queue-draining procedures, managed objects and extensions to existing protocols that enable bridges and end stations to schedule the transmission of frames based on timing derived from IEEE Std 802.1AS.
TSN TAS can be configured at Configuration→TSN→TAS.

TAS Status Parameters

Auto-refresh ☐ Refresh

Port	Oper Gate								Cycle Time			Time		Config Change		Tick Granularity	Config Pending	Gate Control List	
	Enabled	Q0	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Value	Unit	Extension, ns	Base	Current	Time	Error			Length	GCL
1	✗	✓	✓	✓	✓	✓	✓	✓	✓	100	MilliSeconds	256	0	1228312	0	0	1	false	0	Status
2	✗	✓	✓	✓	✓	✓	✓	✓	✓	100	MilliSeconds	256	0	1228312	0	0	1	false	0	Status
3	✗	✓	✓	✓	✓	✓	✓	✓	✓	100	MilliSeconds	256	0	1228312	0	0	1	false	0	Status
4	✗	✓	✓	✓	✓	✓	✓	✓	✓	100	MilliSeconds	256	0	1228312	0	0	1	false	0	Status
5	✗	✓	✓	✓	✓	✓	✓	✓	✓	100	MilliSeconds	256	0	1228312	0	0	1	false	0	Status
6	✗	✓	✓	✓	✓	✓	✓	✓	✓	100	MilliSeconds	256	0	1228312	0	0	1	false	0	Status
7	✗	✓	✓	✓	✓	✓	✓	✓	✓	100	MilliSeconds	256	0	1228312	0	0	1	false	0	Status
8	✗	✓	✓	✓	✓	✓	✓	✓	✓	100	MilliSeconds	256	0	1228312	0	0	1	false	0	Status
9	✗	✓	✓	✓	✓	✓	✓	✓	✓	100	MilliSeconds	256	0	1228312	0	0	1	false	0	Status
10	✗	✓	✓	✓	✓	✓	✓	✓	✓	100	MilliSeconds	256	0	1228312	0	0	1	false	0	Status
11	✗	✓	✓	✓	✓	✓	✓	✓	✓	100	MilliSeconds	256	0	1228312	0	0	1	false	0	Status
12	✗	✓	✓	✓	✓	✓	✓	✓	✓	100	MilliSeconds	256	0	1228312	0	0	1	false	0	Status

Parameter	Description
Port	Port number of the switch.
Oper Gate Enabled	The Enabled parameter shows whether traffic scheduling is active (true) or inactive (false).
Oper Gate States	The current state of the gate associated with each queue for the Port.
Cycle Time Value	The operational value of the gating cycle for the Port. The Cycle Time variable is a rational number of seconds, defined by value and a unit.
Cycle Time Unit	The operational Cycle Time unit. May be milliseconds, microseconds or nanoseconds.
Cycle Time Extension	An integer number of nanoseconds, defining the maximum amount of time by which the gating cycle for the Port is permitted to be extended when a new cycle configuration is installed.
Base Time	The operational value of base time, expressed as an IEEE 1588 precision time protocol (PTP) timescale.
Current Time	The current time, in PTPtime, as maintained by the local system. The value is a representation of a PTPtime value, consisting of a 48-bit integer number of seconds and a 32-bit integer number of nanoseconds. Only the seconds are displayed.
Config Change Time	The PTPtime at which the next config change is scheduled to occur. The value is a representation of a PTPtime value, consisting of a 48-bit integer number of seconds and a 32-bit integer number of nanoseconds.
Config Change Error	A counter of the number of times that a re-configuration of the traffic schedule has been requested with the old schedule still running and the requested base time was in the past.
Tick Granularity	The granularity of the cycle time clock, represented as an unsigned number of tenths of nanoseconds.
Config Pending	The value of the ConfigPending state machine variable. The value is TRUE if a configuration change is in progress but has not yet completed.
GCL Length	The operational value of the gate control list length parameter for the Port. The integer value indicates the number of entries (TLVs) in the operational gate control list.
GCL	A link to the GCL parameter status.

GCL Operational Parameters

This shows the GCL Opertional Parameters status.

GCL Operational Parameters

GCL ID	Gate State								Time Interval
	Q0	Q1	Q2	Q3	Q4	Q5	Q6	Q7	nanoseconds
No entry exists									

Parameter	Description
GCL ID	Index of Gate control list.

Gate State	The GateState shows for each queue whether it is open or closed in the given time interval.
Time Interval	TimeInterval is specified in number of nanoseconds.

5.27.3 PSFP

PSFP (Per Stream Filtering and Policing) functions allow filtering and policing decisions, and subsequent frame queuing decisions on a per-stream basis. PSFP is supported by a table of stream filters that determine the filtering and policing actions that are to be applied to frames received on ingress ports.

5.27.3.1 Capabilities

This shows user how to check the PSFP parameters supported by this platform.
TSN PSFP Capabilities can be configured at Configuration→TSN→PSFP.

PSFP Capabilities

Flow Meter Instance Count	255
Stream Gate Instance Count	255
Stream Gate Control List Length	4
Stream Filter Instance Count	255

Parameter	Description
Flow Meter Instance Count	The number of flow meters that can be instantiated. They are numbered [0; Max].
Stream Gate Instance Count	The number of stream gates that can be instantiated. They are numbered [0; Max].
Stream Gate Control List Length	The maximum number of gate control list entries per stream gate instance.
Stream Filter Instance Count	The number of stream filters that can be instantiated. They are numbered [0; Max].

5.27.3.2 Stream Gate Status

This shows an overview of the status of all PSFP stream gate instances.
TSN PSFP Stream Gate can be configured at Configuration→TSN→PSFP→Stream Gates.

PSFP Stream Gate Status Overview

Stream Gate ID	Enabled	Config Pending	State	IPV	Closed Due To	
					Invalid Rx	Octects Exceeded
No Stream Gate instances						

Parameter	Description
Stream Gate ID	ID of the stream gate instance.
Enabled	This shows whether the gate is administratively enabled (Yes) or disabled (No). The remaining fields are only valid when the gate is enabled.

Config Pending	A stream gate can be configured to take on a set of parameters at a given time in the future. If that time has not been reached, this field will show the value "Yes", otherwise "No". Only filled in if the stream gate is enabled.
State	The operational value of the stream gate's gate state. It reads "Open" if the gate is open, "Closed" if closed. Only filled in if the stream gate is enabled.
IPV	The operational value of the Internal Priority Value for the gate. If "Disabled", use of IPV is not in effect. Otherwise, it contains a value between 0 and 7, which in turn is used to direct the frame to a particular egress queue. If gate control lists are used, frames may later on get directed to other egress queues. Only filled in if the stream gate is enabled.
Closed Due To Invalid Rx	A stream gate may get permanently closed if receiving a frame during a closed gate state. This field reads "Yes" if this is the case, "No" if not. Click the "Clear" button to re-enable the gate. This button is only available if the gate is operationally enabled and is actually in this state. The field as such is only filled in if the stream gate is enabled.
Closed Due To Octets Exceeded	A stream gate may get permanently closed if receiving a frame that exceeds the configured Octet Max and the feature is enabled. This field reads "Yes" if this is the case, "No" if not. Click the "Clear" button to re-enable the gate. This button is only available if the gate is operationally enabled and is actually in this state. The field as such is only filled in if the stream gate is enabled.

5.27.3.3 Stream Filter Status

This shows the current PSFP stream filter status.
TSN PSFP Stream Filter can be configured at Configuration→TSN→PSFP→Stream Filters.

PSFP Stream Filter Status

Clear	Stream Filter ID	Blocked due to Oversize Frame	Warnings
No PSFP Stream Filter instances			

Parameter	Description
Clear	Click this button to clear the blocked flag for this stream filter instance.
Stream Filter ID	The ID of the stream filter instance.
Blocked due to Oversize Frame	If "Yes", the stream is blocked and all frames are discarded. Click the "Clear" or "Clear All" button to clear this flag. If "No", the stream is not blocked.
Warnings	Configuration of a stream filter may result in configurational warnings. See the Configuration→TSN→PSFP→Stream Filters page for details on this.

5.27.3.4 Stream Filter Statistics

This shows the current PSFP stream filter statistics.
TSN PSFP Stream Filter can be configured at Configuration→TSN→PSFP→Stream Filters.

PSFP Stream Filter Statistics

Auto-refresh ☐ Refresh Clear All

Clear	Stream Filter ID	Matching	Passing	Not Passing	Passing SDU	Not Passing SDU	Red
No PSFP Stream Filter instances							

Parameter	Description
Clear	Click this button to clear the statistics for this stream filter instance.
Stream Filter ID	The ID of the stream filter instance.
Matching	Counts received frames that match this stream filter.
Passing	Counts received frames that pass the gate associated with this stream filter.

Not Passing	Counts received frames that do not pass the gate associated with this stream filter.
Passing SDU	Counts received frames that pass the SDU size filter specification associated with this stream filter.
Not Passing SDU	Counts received frames that do not pass the SDU size filter specification associated with this stream filter.
Red	Counts received frames that were discarded as a result of the flow meter associated with this stream filter.

5.27.4 FRER

FRER is an acronym for **F**rame **R**eplication and **E**limination for **R**eliability. Frame Replication and Elimination for Reliability is specified by the IEEE 802.1CB-2017 standard and provides increased reliability (reduced packet loss rates) for a stream.
This is achieved by:

1. Sequence numbering and replicating packets in the talker (source) end system and/or in relay systems in the network.
2. Eliminating those replicates in the listener (destination) end system and/or in other relay systems.

5.27.4.1 Status

This shows users how to check the current FRER status.
TSN FRER can be configured at Configuration→TSN→FRER.

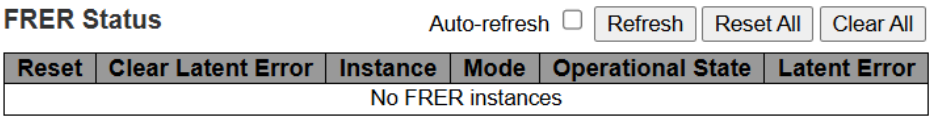

Parameter	Description
Reset	This button is only available on active FRER instances. It allows for resetting the FRER instance. In Generation mode, a click will cause the sequence number generator to start over. In Recovery mode, a click will reset individual recovery functions (if applicable) and the compound recovery functions.
Clear Latent Error	This button is only available on enabled FRER instances in recovery mode with latent error detection enabled. Click to clear a sticky latent error.
Instance	The ID of the FRER instance.
Mode	Shows the mode (Generation/Recovery) of the FRER instance.
Operational State	This shows the current operational state of the selected FRER instance. See the Configuration→TSN→FRER page for details on image color and possible warnings.
Latent Error	This shows the latent error status. See the Configuration→TSN→FRER page for details on image color.
Reset All	Corresponds to clicking on each of the FRER instances' Reset button.
Clear All	Corresponds to clicking on each of the FRER instances' Clear button.

5.27.4.2 Statistics

This shows users how to check the current FRER statistics counters.
TSN FRER can be configured at Configuration→TSN→FRER.

FRER Statistics

Auto-refresh ☐ Refresh Clear All

Clear	Instance	Mode	Egress Port	Ingress Stream	Out of Order	Rogue	Passed	Discarded	Lost	Tagless	Latent Error Resets	Generation Matches	Resets
No FRER instances													

Parameter	Description
Clear	Press this button to clear the counters for the particular instance. In recovery mode, not all rows may have a clear button. That's because a single instance may have several sets of statistics and therefore span several rows.
Instance	The FRER instance ID. In recovery mode, not all rows may have filled out the Instance. That's because a single instance may have several sets of statistics and therefore span several rows. For an empty instance, the instance number is the closest previous row in the table with a non-empty instance number.
Mode	Mode of operation. Generation or Recovery. In recovery mode, not all rows may have filled out the Mode. That's because a single instance may have several sets of statistics and therefore span several rows. For an empty mode, the mode is the closest previous row in the table with a non-empty mode.
Egress Port	Only valid in recovery mode. In recovery mode, there is a set of statistics per egress port. This shows the egress port - both interface name and port number - for which this row of statistics pertains. If empty, it is the same as the egress port presented in the closest previous non-empty row.
Ingress Stream	Only valid in recovery mode. In individual recovery mode, there is a set of statistics per Stream ID as well as a compound set of statistics. In non-individual recovery mode, there is only a compound set of statistics.
Out of Order	IEEE 802.1CB-2017: frerCpsSeqRcvyOutOfOrderPackets. Only valid in Recovery mode.
Rogue	IEEE 802.1CB-2017: frerCpsSeqRcvyRoguePackets. Only valid in Recovery mode.
Passed	IEEE 802.1CB-2017: frerCpsSeqRcvyPassedPackets. Only valid in Recovery mode.
Discarded	IEEE 802.1CB-2017: frerCpsSeqRcvyDiscardedPackets. Only valid in Recovery mode.
Lost	IEEE 802.1CB-2017: frerCpsSeqRcvyLostPackets. Only valid in Recovery mode.
Tagless	IEEE 802.1CB-2017: frerCpsSeqRcvyTaglessPackets. Only valid in Recovery mode.
Latent Error Resets	Only valid in recovery mode and only counts if Latent Error Detection is enabled. Counts the number of times the latent error reset function is invoked. Corresponds to IEEE 802.1CB-2017's frerCpsSeqRcvyLatentErrorResets. Only the first row of a sequence of rows that pertain to this FRER instance is filled in with its value, because it is identical for all statistics sets for a given FRER instance.
Generation Matches	Number of matches on ingress stream. Only valid in Generation mode.
Resets	In generation mode, this corresponds to IEEE 802.1CB-2017's frerCpsSeqGenResets In recovery mode, this corresponds to IEEE 802.1CB-2017's frerCpsSeqRcvyResets.
Clear All	Clears the counters for all entries

5.28 OSPF

OSPF is an acronym for **O**pen **S**hortest **P**ath **F**irst. OSPF is a link-state interior gateway protocol used within a single autonomous system. Each router builds a link-state database of the network topology and applies the SPF algorithm to determine the shortest paths to all destination networks.

5.28.1 Status

This shows OSPF router status table. It is used to provide the OSPF router status information.
OSPF can be configured at Configuration→OSPF.

OSPF Global Status Clear OSPF Process

Status Information	
Router ID	192.168.203.160
SPF Delay	200 msecs
SPF Hold Time	400 msecs
SPF Max. Wait Time	10000 msecs
Last Executed SPF Time Stamp	0 msecs
Min. LSA Interval	5 secs
Min. LSA Arrival	1000 msecs
External LSA Count	0 msecs
External LSA Checksum	0x0
Attached Area Count	0

Parameter	Description
Remote ID	OSPF router ID
SPF Delay	Delay time (in seconds) of SPF calculations.
SPF Hold Time	Minimum hold time (in milliseconds) between consecutive SPF calculations.
SPF Max. Wait Time	Maximum wait time (in milliseconds) between consecutive SPF calculations.
Last Executed SPF Time Stamp	Time (in milliseconds) that has passed between the start of the SPF algorithm execution and the current time.
Min. LSA Interval	Minimum interval (in seconds) between link-state advertisements.
Min. LSA Arrival	Maximum arrival time (in milliseconds) of link-state advertisements.
External LSA Count	Number of external link-state advertisements.
External LSA Checksum	Number of external link-state checksum.
Attached Area Count	Number of areas attached for the router.
Clear OSPF Process	Click to reset the current OSPF process.

5.28.2 Area

This shows OSFP network area status table. It is used to provide the OSPF network area status information. OSPF can be configured at Configuration→OSPF.

OSPF Area Status

Area ID	Backbone	Area Type	NSSA translator State	Active Interfaces	Auth. Type	SPF Executed Times	LSA Count	Router LSA		Network LSA		Summary LSA		ASBR Summary LSA		NSSA LSA	
								Count	Checksum	Count	Checksum	Count	Checksum	Count	Checksum	Count	Checksum
No entry exists																	

Parameter	Description
Area ID	The Area ID.
Backbone	Indicate if it's backbone area or not.
Area Type	The area type.
NSSA translator state	Indicate the current state of the NSSA-ABR translator which the router uses to translate Type-7 LSAs in the NSSA to Type-5 LSAs in backbone area.
Active Interfaces	Number of active interfaces attached in the area.

Auth. Type	The authentication type in the area.
SPF Executed Times	Number of times SPF algorithm has been executed for the particular area.
LSA Count	Number of the total LSAs for the particular area.
Router LSA Count	Number of the router-LSAs (Type-1) of a given type for the particular area.
Router LSA Checksum	The router-LSAs (Type-1) checksum.
Network LSA Count	Number of the network-LSAs (Type-2) of a given type for the particular area.
Network LSA Checksum	The network-LSAs (Type-2) checksum.
Summary LSA Count	Number of the summary-LSAs (Type-3) of a given type for the particular area.
Summary LSA Checksum	The summary-LSAs (Type-3) checksum.
ASBR Summary LSA Count	Number of the ASBR-summary-LSAs (Type-4) of a given type for the particular area.
ASBR Summary LSA Checksum	The ASBR-summary-LSAs (Type-4) checksum.
NSSA LSA Count	Number of the NSSA LSAs of a given type for the particular area.
NSSA LSA Checksum	The NSSA LSAs checksum.

5.28.3 Neighbor

This shows OSPF IPv4 neighbor status table. It is used to provide OSPF neighbor status information.
OSPF can be configured at Configuration→OSPF.

OSPF Neighbor Status

Neighbor ID	Priority	State	Dead Time	Interface Address	Interface
No entry exists					

Parameter	Description
Neighbor ID	The Neighbor ID.
Priority	The priority of OSPF neighbor. It indicates the priority of the neighbor router. This item is used when selecting the DR for the network. The router with the highest priority becomes the DR.
State	The state of OSPF neighbor. It indicates the functional state of the neighbor router.
Dead Time	Dead timer. It indicates the amount of time remaining that the router waits to receive an OSPF hello packet from the neighbor before declaring the neighbor down.
Interface Address	The IP address.
Interface	The network interface.

5.28.4 Interface

This shows OSPF interface status table. It is used to provide OSPF interface status information.
OSPF can be configured at Configuration→OSPF.

OSPF Interface Status

Interface	Interface Address	Area ID	Router ID	State	DR		BDR		Pri	Cost	Interval Configuration(sec)				Hello Timer	Nbr Count	Adjacent Nbr Count	Passive	Transmit Delay
					ID	Address	ID	Address			Hello	Dead	Wait	Retransmit					
No entry exists																			

Parameter	Description
Interface	Interface identification.
Interface Address	IPv4 network address.
Area ID	The OSPF area ID.
Router ID	The OSPF router ID.
State	The state of the link.
DR ID	The router ID of DR.
DR Address	The IP address of DR.
BDR ID	The router ID of BDR.
BDR Address	The IP address of BDR.
Priority	The OSPF priority. It helps determine the DR and BDR on the network to which this interface is connected.
Cost	The cost of the interface.
Hello	Hello timer. A time interval that a router sends an OSPF hello packet.
Dead	Dead timer. Dead timer is a time interval to wait before declaring a neighbor dead. The unit of time is the second.
Wait	This interval is used in Wait Timer. Wait timer is a single shot timer that causes the interface to exit waiting and select a DR on the network. Wait Time interval is the same as Dead time interval.
Retransmit	Retransmit timer. A time interval to wait before retransmitting a database description packet when it has not been acknowledged.
Hello Timer	Hello due timer. An OSPF hello packet will be sent on this interface after this due time.
Nbr Count	Neighbor count. This is the number of OSPF neighbors discovered on this interface.
Adjacent Nbr Count	Adjacent neighbor count. This is the number of routers running OSPF that are fully adjacent with this router.
Passive	Indicate if the interface is passive interface.
Transmit Delay	The estimated time to transmit a link-state update packet on the interface.

5.28.5 Routing

This shows OSPF routing status table. It is used to provide OSPF routing status information.
OSPF can be configured at Configuration→OSPF.

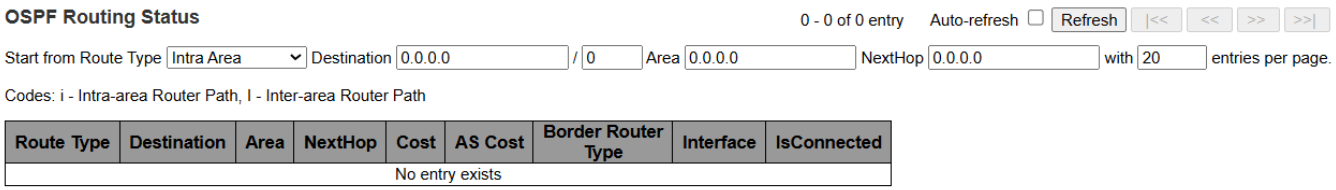

Parameter	Description
Route Type	The OSPF route type. **Intra Area**: The destination is an OSPF route which is located on intra-area. **Inter Area**: The destination is an OSPF route which is located on inter-area. **Border Router**: The destination is a border router. **External Type-1**: The destination is an external Type-1 route. **External Type-2**: The destination is an external Type-2 route.
Destination	Network and prefix (example 10.0.0.0/16) of the given route entry.
Area	It indicates which area the route or router can be reached via/to.
NextHop	Ipv4 address encoded as "a.b.c.d", where a-d is a base-10 human readable integer in the range [0-255]
Cost	The cost of the route.
AS Cost	The cost of the route within the OSPF network. It is valid for external Type-2 route and always '0' for other route type.
Border Router Type	The border router type of the OSPF route entry. **i-ABR**: The border router is an ABR. **i-ASBR**: The border router is an ASBR located on Intra-area. **i-ASBR**: The border router is an ASBR located on Inter-area. **i-ABR/ASBR**: The border router is an ASBR attached to at least 2 areas.
Interface	The interface where the ip packet is outgoing.
IsConnected	The destination is connected directly or not.

5.28.6 Database

In OSPF, each router maintains a separate link-state database (LSDB) for every area it participates in. The LSDB stores all link-state advertisements (LSAs) describing the topology of that area. Neighbors synchronize their databases during adjacency formation by exchanging LSA headers and, when required, the full LSA contents. This ensures that all routers in the same area share a consistent view of the network, which is then used by the SPF algorithm to calculate shortest-path routes

5.28.6.1 General Database

This shows OSFP link state database information table.
OSPF can be configured at Configuration→OSPF.

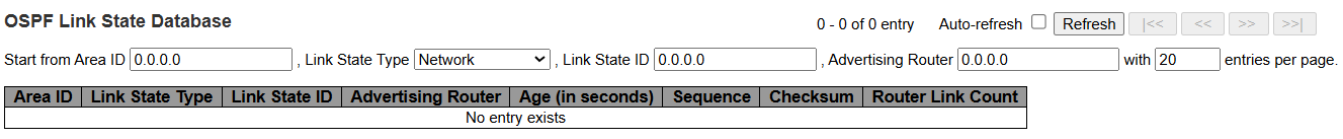

Parameter	Description
Area ID	The OSPF area ID of the link state advertisement. It is not required for external LSA.

Link State Type	The type of the link state advertisement.
Link State ID	The OSPF link state ID. It identifies the piece of the routing domain that is being described by the LSA.
Advertising Router	The advertising router ID which originated the LSA.
Age	The time in seconds since the LSA was originated.
Sequence	The LS sequence number of the LSA.
Checksum	The checksum of the LSA contents.
Router Link Count	The link count of the LSA. The field is significant only when the link state type is 'Router Link State' (Type 1).

5.28.6.2 Detail Database

Router

This shows OSPF LSA Router link state database information table.
OSPF can be configured at Configuration→OSPF.

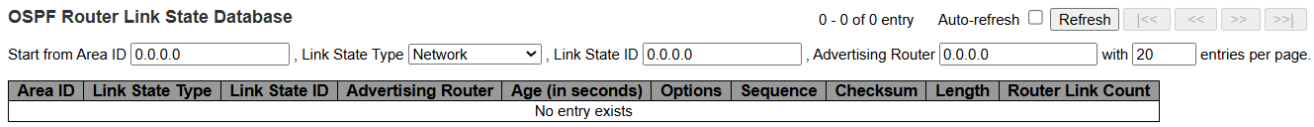

Parameter	Description
Area ID	The OSPF area ID of the link state advertisement
Link State Type	The type of the link state advertisement.
Link State ID	The OSPF link state ID. It identifies the piece of the routing domain that is being described by the LSA.
Advertising Router	The advertising router ID which originated the LSA.
Age	The time in seconds since the LSA was originated.
Options	The OSPF option field which is present in OSPF hello packets, which enables OSPF routers to support (or not support) optional capabilities, and to communicate their capability level to other OSPF routers.
Sequence	The LS sequence number of the LSA.
Checksum	The checksum of the LSA contents.
Length	The Length in bytes of the LSA.
Router Link Count	The link count of the LSA. The field is significant only when the link state type is 'Router Link State' (Type 1).

Network

This shows OSPF LSA network link state database information table.
OSPF can be configured at Configuration→OSPF.

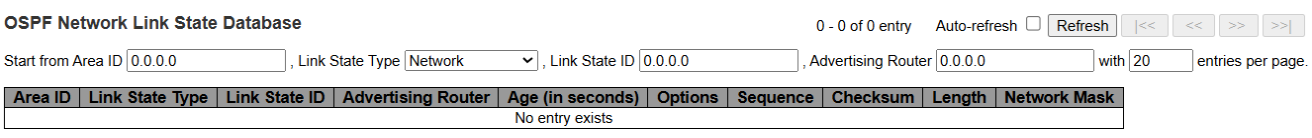

Parameter	Description
Area ID	The OSPF area ID of the link state advertisement
Link State Type	The type of the link state advertisement.
Link State ID	The OSPF link state ID. It identifies the piece of the routing domain that is being described by the LSA.
Advertising Router	The advertising router ID which originated the LSA.
Age	The time in seconds since the LSA was originated.
Options	The OSPF option field which is present in OSPF hello packets, which enables OSPF routers to support (or not support) optional capabilities, and to communicate their capability level to other OSPF routers.
Sequence	The LS sequence number of the LSA.
Checksum	The checksum of the LSA contents.
Length	The Length in bytes of the LSA.
NetworkMask	Network mask length. The field is significant only when the link state type is 'Network Link State' (Type 2).

Summary

This shows OSPF LSA Summary link state database information table.
OSPF can be configured at Configuration→OSPF.

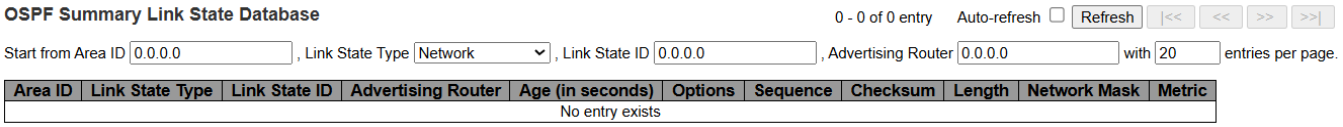

Parameter	Description
Area ID	The OSPF area ID of the link state advertisement
Link State Type	The type of the link state advertisement.
Link State ID	The OSPF link state ID. It identifies the piece of the routing domain that is being described by the LSA.
Advertising Router	The advertising router ID which originated the LSA.
Age	The time in seconds since the LSA was originated.
Options	The OSPF option field which is present in OSPF hello packets, which enables OSPF routers to support (or not support) optional capabilities, and to communicate their capability level to other OSPF routers.
Sequence	The LS sequence number of the LSA.
Checksum	The checksum of the LSA contents.
Length	The Length in bytes of the LSA.
NetworkMask	Network mask length. The field is significant only when the link state type is 'Summary/ASBR Summary Link State' (Type 3, 4).
Metric	User specified metric for this summary route. The field is significant only when the link state type is 'Summary/ASBR Summary Link State' (Type 3, 4).

ASBR Summary

This shows OSPF LSA Summary link state database information table.
OSPF can be configured at Configuration→OSPF.

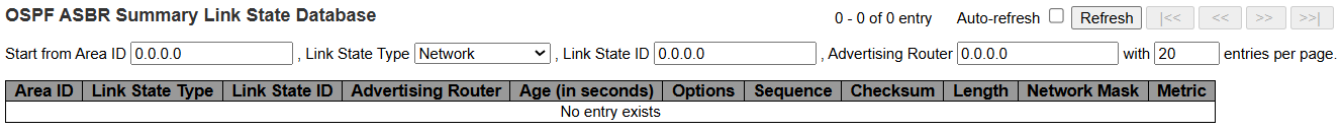

Parameter	Description
Area ID	The OSPF area ID of the link state advertisement
Link State Type	The type of the link state advertisement.
Link State ID	The OSPF link state ID. It identifies the piece of the routing domain that is being described by the LSA.
Advertising Router	The advertising router ID which originated the LSA.
Age	The time in seconds since the LSA was originated.
Options	The OSPF option field which is present in OSPF hello packets, which enables OSPF routers to support (or not support) optional capabilities, and to communicate their capability level to other OSPF routers.
Sequence	The LS sequence number of the LSA.
Checksum	The checksum of the LSA contents.
Length	The Length in bytes of the LSA.
NetworkMask	Network mask length. The field is significant only when the link state type is 'Summary/ASBR Summary Link State' (Type 3, 4).
Metric	User specified metric for this summary route. The field is significant only when the link state type is 'Summary/ASBR Summary Link State' (Type 3, 4).

External

This shows OSPF LSA External link state database information table.
OSPF can be configured at Configuration→OSPF.

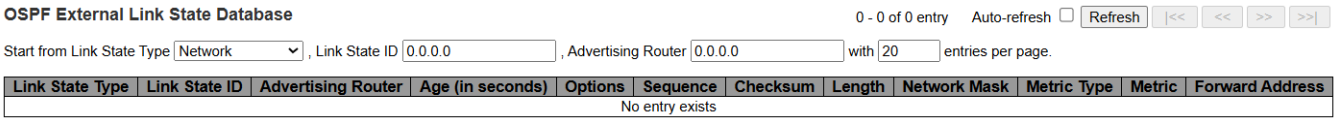

Parameter	Description
Link State Type	The type of the link state advertisement.
Link State ID	The OSPF link state ID. It identifies the piece of the routing domain that is being described by the LSA.
Advertising Router	The advertising router ID which originated the LSA.
Age	The time in seconds since the LSA was originated.
Options	The OSPF option field which is present in OSPF hello packets, which enables OSPF routers to support (or not support) optional capabilities, and to communicate their capability level to other OSPF routers.
Sequence	The LS sequence number of the LSA.

Checksum	The checksum of the LSA contents.
Length	The Length in bytes of the LSA.
NetworkMask	Network mask length. The field is significant only when the link state type is 'External/NSSA External Link State' (Type 5, 7).
Metric Type	The External type of the LSA. The field is significant only when the link state type is 'External/NSSA External Link State' (Type 5, 7).
Metric	User specified metric for this summary route. The field is significant only when the link state type is 'External/NSSA External Link State' (Type 5, 7).
Forward Address	The IP address of forward address. The field is significant only when the link state type is 'External/NSSA External Link State' (Type 5, 7).

NSSA External

This shows OSPF LSA NSSA External link state database information table.
OSPF can be configured at Configuration→OSPF.

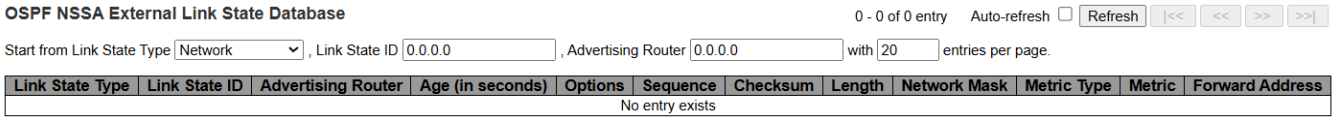

Parameter	Description
Link State Type	The type of the link state advertisement.
Link State ID	The OSPF link state ID. It identifies the piece of the routing domain that is being described by the LSA.
Advertising Router	The advertising router ID which originated the LSA.
Age	The time in seconds since the LSA was originated.
Options	The OSPF option field which is present in OSPF hello packets, which enables OSPF routers to support (or not support) optional capabilities, and to communicate their capability level to other OSPF routers.
Sequence	The LS sequence number of the LSA.
Checksum	The checksum of the LSA contents.
Length	The Length in bytes of the LSA.
NetworkMask	Network mask length. The field is significant only when the link state type is 'External/NSSA External Link State' (Type 5, 7).
Metric Type	The External type of the LSA. The field is significant only when the link state type is 'External/NSSA External Link State' (Type 5, 7).
Metric	User specified metric for this summary route. The field is significant only when the link state type is 'External/NSSA External Link State' (Type 5, 7).
Forward Address	The IP address of forward address. The field is significant only when the link state type is 'External/NSSA External Link State' (Type 5, 7).

5.29 OSPFv3

OSPFv3 is an acronym for **O**pen **S**hortest **P**ath First **v**ersion **3**. OSPFv3 is a link-state interior gateway protocol for IPv6 used within a single autonomous system. Each router builds an IPv6 link-state database of the network topology and applies the SPF algorithm to determine the shortest paths to all IPv6 destination networks.

5.29.1 Status

This shows OSPF6 router status table. It is used to provide the OSPF6 router status information. OSPFv3 can be configured at Configuration→OSPFv3.

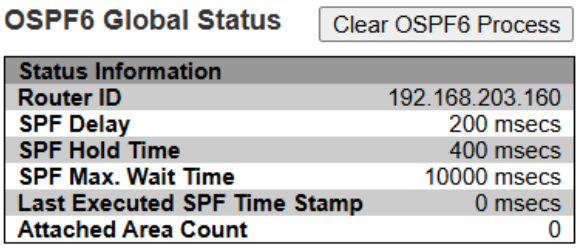

Parameter	Description
Remote ID	OSPF6 router ID
SPF Delay	Delay time (in seconds) of SPF calculations.
SPF Hold Time	Minimum hold time (in milliseconds) between consecutive SPF calculations.
SPF Max. Wait Time	Maximum wait time (in milliseconds) between consecutive SPF calculations.
Last Executed SPF Time Stamp	Time (in milliseconds) that has passed between the start of the SPF algorithm execution and the current time.
Attached Area Count	Number of areas attached for the router.
Clear OSPF6 Process	Click to reset the current OSPF6 process.

5.29.2 Area

This shows OSPF6 network area status table. It is used to provide the OSPF6 network area status information. OSPFv3 can be configured at Configuration→OSPFv3.

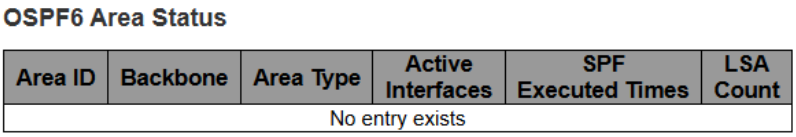

Parameter	Description
Area ID	The Area ID.
Backbone	Indicate if it's backbone area or not.
Area Type	The area type.
Active Interfaces	Number of active interfaces attached in the area.
SPF Executed Times	Number of times SPF algorithm has been executed for the particular area.
LSA Count	Number of the total LSAs for the particular area.

5.29.3 Neighbor

This shows OSPF6 IPv6 neighbor status table. It is used to provide OSPF6 neighbor status information. OSPFv3 can be configured at Configuration→OSPFv3.

OSPF6 Neighbor Status

Neighbor ID	Priority	State	Dead Time	Interface Address	Interface
No entry exists					

Parameter	Description
Neighbor ID	The Neighbor ID.
Priority	The priority of OSPF6 neighbor. It indicates the priority of the neighbor router. This item is used when selecting the DR for the network. The router with the highest priority becomes the DR.
State	The state of OSPF6 neighbor. It indicates the functional state of the neighbor router.
Dead Time	Dead timer. It indicates the amount of time remaining that the router waits to receive an OSPF6 hello packet from the neighbor before declaring the neighbor down.
Interface Address	The IP address.
Interface	The network interface.

5.29.4 Interface

This shows OSPF6 interface status table. It is used to provide OSPF6 interface status information. OSPFv3 can be configured at Configuration→OSPFv3.

OSPF6 Interface Status

Interface	Interface Address	Area ID	Router ID	State	DR ID	BDR ID	Pri	Cost	Interval Configuration(sec)			Passive	Transmit Delay
									Hello	Dead	Retransmit		
No entry exists													

Parameter	Description
Interface	Interface identification.
Interface Address	IPv6 network address.
Area ID	The OSPF6 area ID.
Router ID	The OSPF6 router ID.
State	The state of the link.
DR ID	The router ID of DR.
DR Address	The IP address of DR.
BDR ID	The router ID of BDR.
BDR Address	The IP address of BDR.
Priority	The OSPF6 priority. It helps determine the DR and BDR on the network to which this interface is connected.

Cost	The cost of the interface.
Hello	Hello timer. A time interval that a router sends an OSPF6 hello packet.
Dead	Dead timer. Dead timer is a time interval to wait before declaring a neighbor dead. The unit of time is the second.
Retransmit	Retransmit timer. A time interval to wait before retransmitting a database description packet when it has not been acknowledged.
Passive	Indicate if the interface is passive interface.
Transmit Delay	The estimated time to transmit a link-state update packet on the interface.

5.29.5 Routing

This shows OSPF6 routing status table. It is used to provide OSPF6 routing status information.
OSPFv3 can be configured at Configuration→OSPFv3.

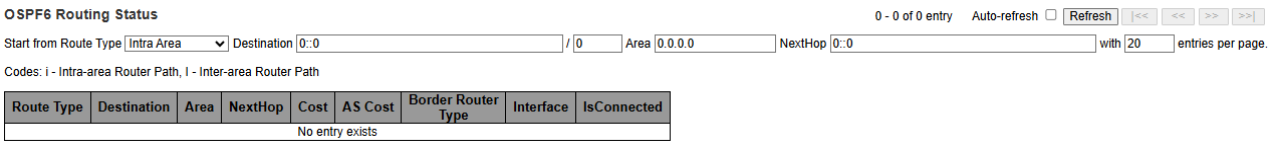

Parameter	Description
Route Type	The OSPF6 route type. **Intra Area**: The destination is an OSPF route which is located on intra-area. **Inter Area**: The destination is an OSPF route which is located on inter-area. **Border Router**: The destination is a border router. **External Type-1**: The destination is an external Type-1 route. **External Type-2**: The destination is an external Type-2 route.
Destination	Network and prefix (example 10.0.0.0/16) of the given route entry.
Area	It indicates which area the route or router can be reached via/to.
NextHop	An Ipv6 address represented as human readable test as specified in RFC5952
Cost	The cost of the route.
AS Cost	The cost of the route within the OSPF6 network. It is valid for external Type-2 route and always '0' for other route type.
Border Router Type	The border router type of the OSPF6 route entry. **i-ABR**: The border router is an ABR. **i-ASBR**: The border router is an ASBR located on Intra-area. **i-ASBR**: The border router is an ASBR located on Inter-area. **i-ABR/ASBR**: The border router is an ASBR attached to at least 2 areas.
Interface	The interface where the ip packet is outgoing.
IsConnected	The destination is connected directly or not.

5.29.6 Database

In OSPFv3, each router maintains a separate link-state database (LSDB) for every IPv6 area it belongs to. The LSDB stores all OSPFv3 link-state advertisements (LSAs), including router, network, inter-area, and external LSAs, that describe the IPv6 topology and reachable prefixes within that area. During adjacency formation, neighbors synchronize their databases by exchanging Database Description (DD) packets containing LSA

headers and then requesting any missing or newer LSAs with Link-State Request packets, receiving the full LSA contents in Link-State Update packets as needed. This process ensures that all routers in the same area share a consistent IPv6 topology view in their LSDBs, which is then used by the SPF algorithm to calculate shortest-path routes and install them into the IPv6 routing table.

5.29.6.1 General Database

This shows OSFP6 link state database information table.
OSPFv3 can be configured at Configuration→OSPFv3

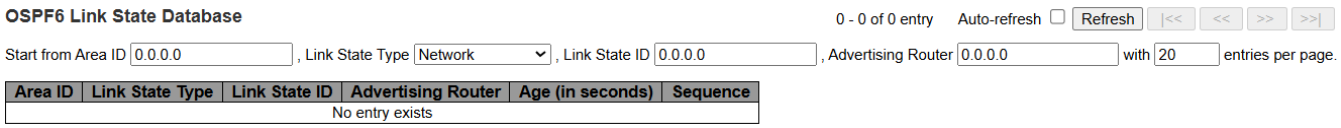

Parameter	Description
Area ID	The OSPF6 area ID of the link state advertisement. It is not required for external LSA.
Link State Type	The type of the link state advertisement.
Link State ID	The OSPF6 link state ID. It identifies the piece of the routing domain that is being described by the LSA.
Advertising Router	The advertising router ID which originated the LSA.
Age	The time in seconds since the LSA was originated.
Sequence	The LS sequence number of the LSA.

5.29.6.2 Detail Database

Router

This shows OSPF6 LSA Router link state database information table.
OSPFv3 can be configured at Configuration→OSPFv3.

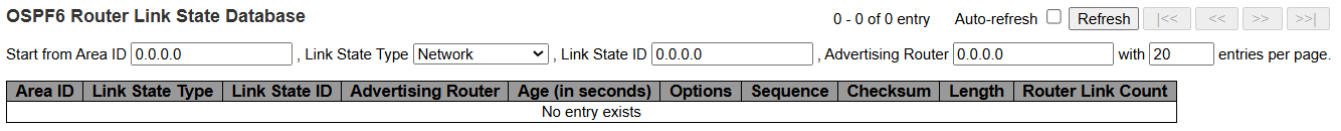

Parameter	Description
Area ID	The OSPF6 area ID of the link state advertisement
Link State Type	The type of the link state advertisement.
Link State ID	The OSPF6 link state ID. It identifies the piece of the routing domain that is being described by the LSA.
Advertising Router	The advertising router ID which originated the LSA.
Age	The time in seconds since the LSA was originated.
Options	The OSPF6 option field which is present in OSPF6 hello packets, which enables OSPF6 routers to support (or not support) optional capabilities, and to communicate their capability level to other OSPF6 routers.
Sequence	The LS sequence number of the LSA.

Checksum	The checksum of the LSA contents.
Length	The Length in bytes of the LSA.
Router Link Count	The link count of the LSA. The field is significant only when the link state type is 'Router Link State' (Type 1).

Network

This shows OSPF6 LSA network link state database information table.
OSPFv3 can be configured at Configuration→OSPFv3.

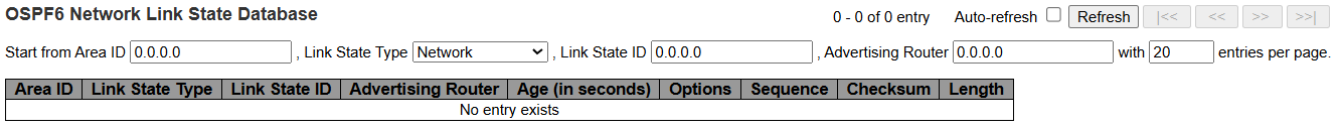

Parameter	Description
Area ID	The OSPF6 area ID of the link state advertisement
Link State Type	The type of the link state advertisement.
Link State ID	The OSPF6 link state ID. It identifies the piece of the routing domain that is being described by the LSA.
Advertising Router	The advertising router ID which originated the LSA.
Age	The time in seconds since the LSA was originated.
Options	The OSPF6 option field which is present in OSPF6 hello packets, which enables OSPF6 routers to support (or not support) optional capabilities, and to communicate their capability level to other OSPF6 routers.
Sequence	The LS sequence number of the LSA.
Checksum	The checksum of the LSA contents.
Length	The Length in bytes of the LSA.

Link

This shows OSPF6 LSA Link state database information table.
OSPFv3 can be configured at Configuration→OSPFv3.

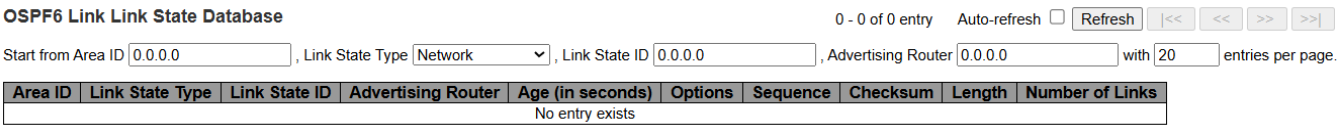

Parameter	Description
Area ID	The OSPF6 area ID of the link state advertisement
Link State Type	The type of the link state advertisement.
Link State ID	The OSPF6 link state ID. It identifies the piece of the routing domain that is being described by the LSA.
Advertising Router	The advertising router ID which originated the LSA.
Age	The time in seconds since the LSA was originated.

Options	The OSPF6 option field which is present in OSPF6 hello packets, which enables OSPF6 routers to support (or not support) optional capabilities, and to communicate their capability level to other OSPF6 routers.
Sequence	The LS sequence number of the LSA.
Checksum	The checksum of the LSA contents.
Length	The Length in bytes of the LSA.
Number of Links	The count of the LSA.

IntraArea Prefix

This shows OSPF6 LSA Intra Area Prefix Link state database information table.
OSPFv3 can be configured at Configuration→OSPFv3.

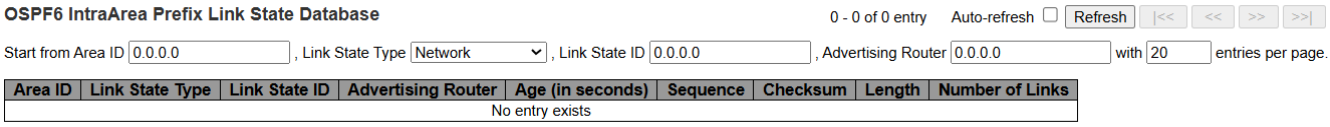

Parameter	Description
Area ID	The OSPF6 area ID of the link state advertisement
Link State Type	The type of the link state advertisement.
Link State ID	The OSPF6 link state ID. It identifies the piece of the routing domain that is being described by the LSA.
Advertising Router	The advertising router ID which originated the LSA.
Age	The time in seconds since the LSA was originated.
Options	The OSPF6 option field which is present in OSPF6 hello packets, which enables OSPF6 routers to support (or not support) optional capabilities, and to communicate their capability level to other OSPF6 routers.
Sequence	The LS sequence number of the LSA.
Checksum	The checksum of the LSA contents.
Length	The Length in bytes of the LSA.
Number of Links	The count of the Prefixes.

InterArea Prefix

This shows OSPF6 LSA Inter Area Prefix Link state database information table.
OSPFv3 can be configured at Configuration→OSPFv3.

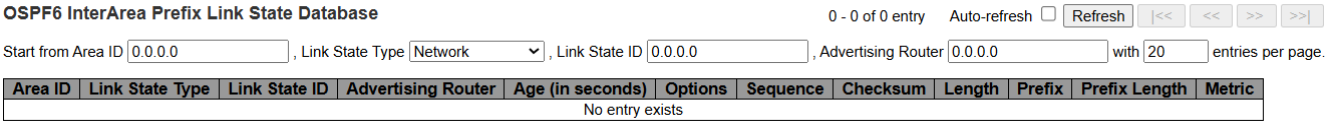

Parameter	Description
Area ID	The OSPF6 area ID of the link state advertisement
Link State Type	The type of the link state advertisement.

Link State ID	The OSPF6 link state ID. It identifies the piece of the routing domain that is being described by the LSA.
Advertising Router	The advertising router ID which originated the LSA.
Age	The time in seconds since the LSA was originated.
Options	The OSPF6 option field which is present in OSPF6 hello packets, which enables OSPF6 routers to support (or not support) optional capabilities, and to communicate their capability level to other OSPF6 routers.
Sequence	The LS sequence number of the LSA.
Checksum	The checksum of the LSA contents.
Length	The Length in bytes of the LSA.
Prefix	IPv6 network address
Prefix Length	IPv6 network mask length.
Metric	User specified metric for this summary route. The field is significant only when the link state type is 'Inter_Area Prefix/Router Link State' (Type 3, 4).

InterArea Router

This shows OSPF6 LSA Inter Area Router Link state database information table.
OSPFv3 can be configured at Configuration→OSPFv3.

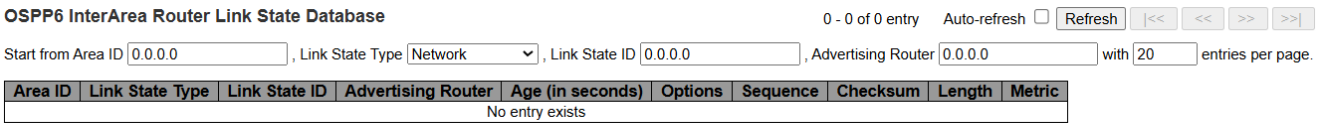

Parameter	Description
Area ID	The OSPF6 area ID of the link state advertisement
Link State Type	The type of the link state advertisement.
Link State ID	The OSPF6 link state ID. It identifies the piece of the routing domain that is being described by the LSA.
Advertising Router	The advertising router ID which originated the LSA.
Age	The time in seconds since the LSA was originated.
Options	The OSPF6 option field which is present in OSPF6 hello packets, which enables OSPF6 routers to support (or not support) optional capabilities, and to communicate their capability level to other OSPF6 routers.
Sequence	The LS sequence number of the LSA.
Checksum	The checksum of the LSA contents.
Length	The Length in bytes of the LSA.
Metric	User specified metric for this summary route. The field is significant only when the link state type is 'Inter_Area Prefix/Router Link State' (Type 3, 4).

External

This shows OSPF6 LSA External Link state database information table.
OSPFv3 can be configured at Configuration→OSPFv3.

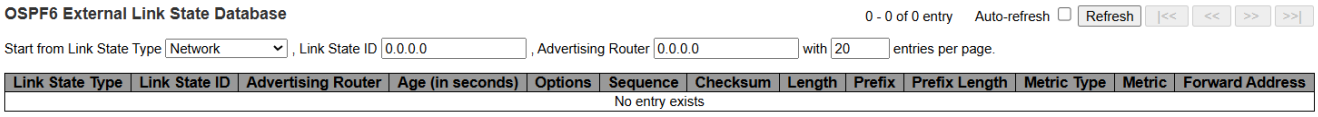

Parameter	Description
Link State Type	The type of the link state advertisement.
Link State ID	The OSPF6 link state ID. It identifies the piece of the routing domain that is being described by the LSA.
Advertising Router	The advertising router ID which originated the LSA.
Age	The time in seconds since the LSA was originated.
Options	The OSPF6 option field which is present in OSPF6 hello packets, which enables OSPF6 routers to support (or not support) optional capabilities, and to communicate their capability level to other OSPF6 routers.
Sequence	The LS sequence number of the LSA.
Checksum	The checksum of the LSA contents.
Length	The Length in bytes of the LSA.
Prefix	IPv6 network address.
Prefix Length	IPv6 network mask length.
Metric Type	The External type of the LSA. The field is significant only when the link state type is 'External/NSSA External Link State' (Type 5, 7).
Metric	User specified metric for this summary route. The field is significant only when the link state type is 'External/NSSA External Link State' (Type 5, 7).
Forward Address	The IP address of forward address. The field is significant only when the link state type is 'External/NSSA External Link State' (Type 5, 7).

5.30 RIP

Routing Information Protocol (RIP) is a distance-vector interior gateway protocol for IPv4 used within a single autonomous system. Each router maintains a routing table with the hop count to each IPv4 destination network and prefers routes with the lowest hop count.

5.30.1 Status

This shows RIP general status information table.
RIP can be configured at Configuration→RIP.

RIP Global Status Clear RIP Process

Status Information	
Version	Default
Update Timer	30 secs
Invalid Timer	180 secs
Garbage-Collection Timer	120 secs
Next Update Time	18 secs
Redistribute Default Metric	1
Redistribute Connected	Disable
Redistribute Static	Disable
Redistribute OSPF	Disable
Administrative Distance	120

Parameter	Description
Version	This indicates the global rip version. By default, the router sends RIPv2 and accepts both RIPv1 and RIPv2. When the router receives either version of REQUESTS or triggered updates packets, it replies with the appropriate version. Be aware that the RIP network class configuration when RIPv1 is involved in the topology. RIPv1 uses classful routing, the subnet information is not included in the routing updates. The limitation makes it impossible to have different-sized subnets inside of the same network class. In other words, all subnets in a network class must have the same size.
Update Timer	The timer interval (in seconds) between the router sends the complete routing table to all neighboring RIP routers
Invalid Timer	The invalid timer is the number of seconds after which a route will be marked invalid.
Garbage-Collection Timer	The garbage collection timer is the number of seconds after which a route will be deleted.
Next Update Time	Specifies when the next round of updates will be sent out from this router in seconds.
Redistribute Default Metric	This indicates the default metric value of redistributed routes.
Redistribute Connected	This indicates the connected route is redistributed or not.
Redistribute Static	This indicates the static route is redistributed or not.
Redistribute OSPF	This indicates the OSPF route is redistributed or not.
Administrative Distance	This indicates administrative distance value
Clear RIP Process	Click to reset the current RIP process.

5.30.2 Interface

This shows RIP interface status information table.
RIP can be configured at Configuration→RIP.

RIP Interface Status

Interface	Send Version	Receive Version	Triggered Update	Passive	Auth. Type	Key-Chain Name
No entry exists						

Parameter	Description
Interface	Interface identification.
Send Version	The RIP version for the advertisement transmission on the interface.

Receive Version	The RIP version for the advertisement reception on the interface.
Triggered Update	This indicates the interface enable triggered update or not.
Passive	This indicates if the passive-interface is active on the interface or not.
Key-Chain Name	This indicates the interface is associate with a specific key-chain name.

5.30.3 Peer

This shows RIP peer status information table.
RIP can be configured at Configuration→RIP.

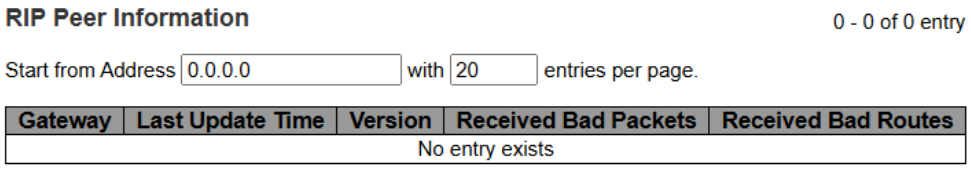

Parameter	Description
Gateway	Peer IPv4 address.
Version	The RIP version number in the header of the last RIP packet received from the neighbor.
Last Update Time	The time duration in seconds from the time the last RIP packet received from the neighbor to now.
Received Bad Packets	The number of RIP response packets from the neighbor discarded as invalid.
Received Bad Routes	The number of routes from the neighbor that were ignored because they were invalid.

5.30.4 Database

This shows RIP database information table.
RIP can be configured at Configuration→RIP.

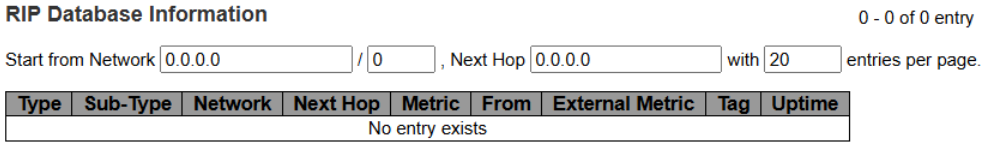

Parameter	Description
Type	The protocol type of the route.
Sub-Type	The protocol sub-type of the route.
Network	The destination IP address and mask of the route.
Next Hop	The first gateway along the route to the destination.
Metric	The metric of the route.

From	This indicates the route is learned an IP address or generated from one of the local interfaces.
External Metric	The field is significant only when the route is redistributed from other protocol type, for example, OSPF. This indicates the metric value from the original redistributed source.
Tag	The tag of the route. It is used to provide a method of separating 'internal' RIP routes, which may have been imported from an EGP (Exterior gateway protocol) or another IGP (Interior gateway protocol). For example, routes imported from OSPF can have a route tag value which the other routing protocols can use to prevent advertising the same route back to the original protocol routing domain.
Uptime	The time field is significant only when the route is learned from the neighbors. When the route destination is reachable (its metric value less than 16), the time field means the invalid time of the route. When the route destination is unreachable (its metric value great than 16), the time field means the garbage-collection time of the route.

5.31 RedBox

RedBox stands for Redundancy Box and covers PRP and HSR functionality. PRP and HSR are specified in IEC 62439-3, ED4(2021) and its CORRIGENDUM 1 (2023).

5.31.1 Status

This shows the status of all created RedBox instances. RedBox can be configured at Configuration→RedBox.

If a RedBox instance is created, but not enabled, a single line spanning all status columns will show 'Inactive'. Otherwise, the columns will be filled with the following information:

RedBox Status

Instance	Mode	Interfaces			Configurational Warnings	Notifications
		Port A	Port B	Port C		
No RedBox instances						

Parameter	Description
Instance	Identifies the RedBox instance number.
Mode	Shows the mode in which the RedBox is currently running. See Configuration→Redbox for a list of possible modes.
Port A	This is the physical interface (port) configured as **L**ink **R**edundancy **E**ntity (LRE) port A. If the RedBox instance connects directly to a neighboring RedBox instance on the same physical device without using physical ports and loop cables, the port is called 'Neighbor'.
Port B	See description of Port A above.
Port C	This is the interlink port. This port number comes indirectly from the configuration of Port A and Port B, as follows: If Port A is 'Neighbor', Port C is the same as Port B. Otherwise Port C is the same as Port A. If both Port A and Port B refer to real ports, Port A is - as said - the interlink port (Port C). Whatever you configure on this port corresponds to configuring both Port A and Port B. Port B is considered unconnected - except for configured port speeds and link.
Configurational Warnings	This shows whether there are configurational warnings that should be solved in order for this RedBox to work properly configuration-wise. Refer to Configuration→Redbox for a description of the colors and a list of possible configurational warnings.
Notifications	At runtime, various conditions that need to attract the operator's attention may arise. Notifications fall into two groups as indicated with a color:

	•  : No runtime errors observed. •  : Runtime errors are observed. If runtime errors are observed, hover the mouse over the image to see a list of the errors. The possible errors are as follows: • NodesTable/ProxyNodeTable is full • Port A has received PRP traffic with wrong Lan ID • Port B has received PRP traffic with wrong Lan ID • Port C has received PRP traffic with wrong Lan ID • Port A has received traffic without an HSR tag • Port B has received traffic without an HSR tag • Port C has received traffic without an HSR tag • Port A's link is down • Port B's link is down A notification disappears when the erroneous condition is no longer detected. It may take up to 10 seconds for the 'Wrong LAN ID' and 'without an HSR tag' conditions to be detected. Once such a condition goes away, it will take at least 20 seconds until the notification disappears.

5.31.2 Statistics

This shows a statistics summary of all created RedBox instances.
Rx and Tx are seen from the RedBox' perspective. 'Port C Rx', for instance, means the number of frames received by the RedBox from the switch core side. If a RedBox instance is created, but not enabled, a single line spanning all status columns will show 'Inactive'.
RedBox can be configured at Configuration→RedBox.

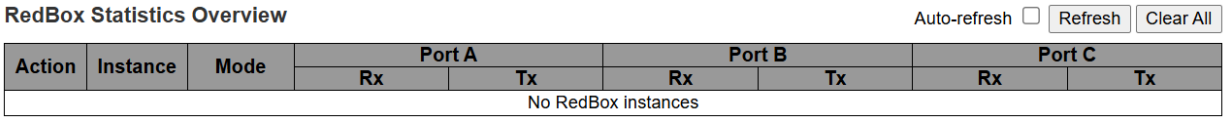

Parameter	Description
Action	This field contains a button that - when clicked - will cause only that entry's statistics to be cleared. The clearing will happen immediately.
Instance	Identifies the RedBox instance number. If active, this also functions as a hyperlink to a detailed statistics page for that instance.
Mode	The mode this RedBox is currently operating in.
Port A/B/C Rx	This indicates the number of frames received on the LRE port. The number includes both HSR-tagged, HSR-untagged, frames with RCT, frames without RCT, as well as number of BPDUs.
Port A/B/C Tx	This indicates the number of frames transmitted on the LRE port. The number includes both HSR-tagged, HSR-untagged, frames with RCT, frames without RCT, as well as number of BPDUs.

5.31.3 Nodes Table

This shows an overview of the contents of the NodesTable of all created RedBox instances.
If a RedBox instance is created, but not enabled, a single line will show 'Inactive'.
RedBox can be configured at Configuration→RedBox.

RedBox Nodes Table

Action	Instance	Mode	MAC Addresses	Wrong LAN
No RedBox instances				

Parameter	Description
Action	This field contains a button that - when clicked - will cause only that entry's NodesTable to be cleared. The clearing will happen immediately.
Instance	Identifies the RedBox instance number. If active, this also functions as a hyperlink to a detailed NodesTable page for that instance.
Mode	The mode this RedBox is currently operating in.
MAC Address	This indicates the number of MAC addresses currently stored in the NodesTable.
Wrong LAN	This is only relevant in PRP-SAN mode. In other modes it is shown with a dash ('-'). If at least one of the entries in the NodesTable has a non-zero Wrong LAN counter, it indicates this with a 'Yes'. Otherwise it reads 'No'.

5.31.4 Proxy Node Table

This page gives an overview of the contents of the ProxyNodeTable of all created RedBox instances.
If a RedBox instance is created, but not enabled, a single line will show 'Inactive'.
RedBox can be configured at Configuration→RedBox.

RedBox ProxyNode Table

Action	Instance	Mode	MAC Addresses	Wrong LAN
No RedBox instances				

Parameter	Description
Action	This field contains a button that - when clicked - will cause only that entry's ProxyNodeTable to be cleared. The clearing will happen immediately. Notice that the RedBox itself inserts two MAC addresses that will never be cleared. One is for the RedBox itself and the other is for the device's management MAC address.
Instance	Identifies the RedBox instance number. If active, this also functions as a hyperlink to a detailed ProxyNodeTable page for that instance.
Mode	The mode this RedBox is currently operating in.
MAC Addresses	This indicates the number of MAC addresses currently stored in the ProxyNodeTable.
Wrong LAN	This is only relevant in HSR-PRP mode. In other modes it is shown with a dash ('-'). If at least one of the entries in the ProxyNodeTable has a non-zero Wrong LAN counter, it indicates this with a 'Yes'. Otherwise it reads 'No'.

6 Diagnostics

Diagnostics is a collection of troubleshooting tools that help verify connectivity, monitor path reachability, and check the physical health of Ethernet links.

6.1 Ping (IPv4)

This page allows you to issue ICMP (IPv4) PING packets to troubleshoot IP connectivity issues.

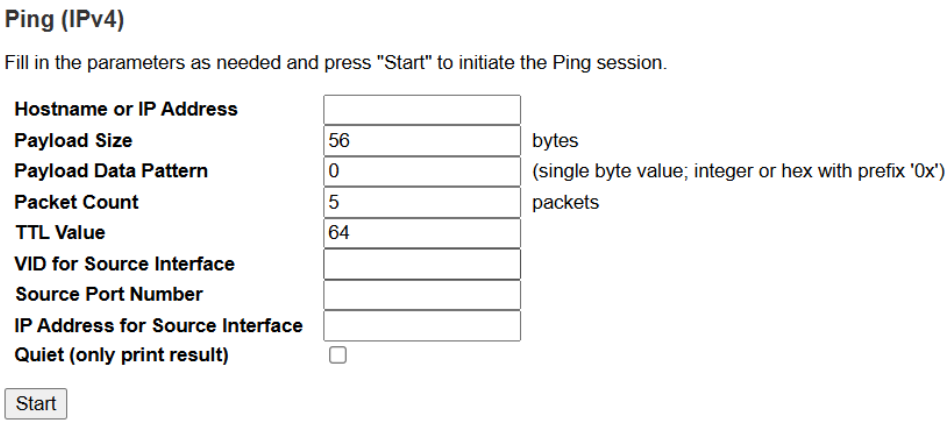

Parameter	Description
Hostname or IP Address	The address of the destination host, either as a symbolic hostname or an IP Address.
Payload Size	Determines the size of the ICMP data payload in bytes (excluding the size of Ethernet, IP and ICMP headers). The default value is 56 bytes. The valid range is **2-1452** bytes.
Payload Data Pattern	Determines the pattern used in the ICMP data payload. The default value is 0. The valid range is **0-255**.
Packet Count	Determines the number of PING requests sent. The default value is 5. The valid range is **1-60**.
TTL Value	Determines the Time-To-Live /TTL) field value in the IPv4 header. The default value is 64. The valid range is **1-255**.
VID for Source Interface	This field can be used to force the test to use a specific local VLAN interface as the source interface. Leave this field empty for automatic selection based on routing configuration. Note: You may only specify either the VID or the IP Address for the source interface.
Source Port Number	This field can be used to force the test to use a specific local interface with the specified port number as the source interface. The specified port must be configured with a suitable IP address. Leave this field empty for automatic selection based on routing configuration. Note: You may only specify either the Source Port Number or the IP Address for the source interface.
Address for Source Interface	This field can be used to force the test to use a specific local interface with the specified IP address as the source interface. The specified IP address must be configured on a local interface. Leave this field empty for automatic selection based on routing configuration. *Note: You may only specify either the VID or the IP Address for the source interface.*
Quiet (only print result)	Checking this option will not print the result of each ping request but will only show the final result. After you press **Start**, ICMP packets are transmitted, and the sequence number and round-trip time are displayed upon reception of a reply. The amount of data received inside of an IP packet of type ICMP ECHO_REPLY will always be 8 bytes more than the requested payload data size (the difference is the ICMP header). The page refreshes automatically until responses to all packets are received, or until a timeout occurs. The output from the command will look like the following: PING 172.16.1.1 (172.16.1.1) from 172.16.1.10: 56 data bytes 64 bytes from 172.16.1.1: seq=0 ttl=64 time=2.034 ms 64 bytes from 172.16.1.1: seq=1 ttl=64 time=1.729 ms 64 bytes from 172.16.1.1: seq=2 ttl=64 time=1.954 ms

	64 bytes from 172.16.1.1: seq=3 ttl=64 time=1.699 ms 64 bytes from 172.16.1.1: seq=4 ttl=64 time=1.916 ms --- 172.16.1.1 ping statistics --- 5 packets transmitted, 5 packets received, 0% packet loss round-trip min/avg/max = 1.699/1.866/2.034 ms

6.2 Ping (IPv6)

This page allows you to issue ICMP (IPv6) PING packets to troubleshoot IPv6 connectivity issues.

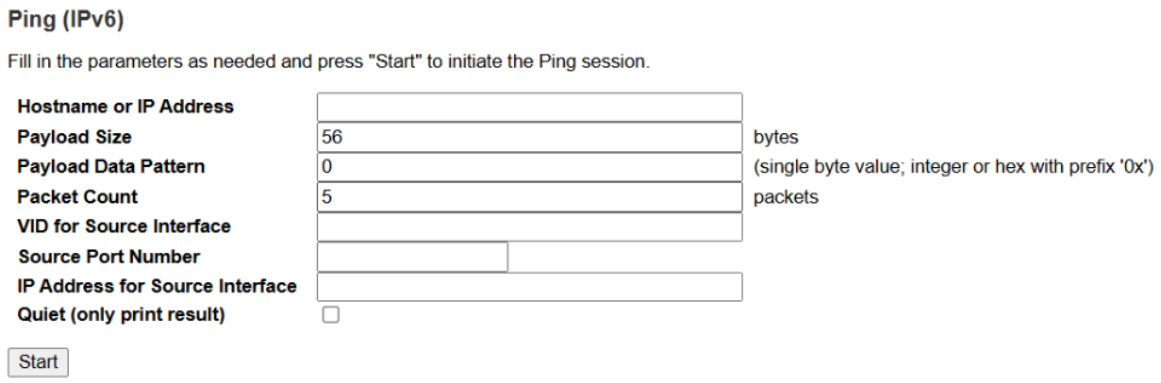

Parameter	Description
Hostname or IP Address	The address of the destination host, either as a symbolic hostname or an IP Address.
Payload Size	Determines the size of the ICMP data payload in bytes (excluding the size of Ethernet, IP and ICMP headers). The default value is 56 bytes. The valid range is **2-1452** bytes.
Payload Data Pattern	Determines the pattern used in the ICMP data payload. The default value is 0. The valid range is **0-255**.
Packet Count	Determines the number of PING requests sent. The default value is 5. The valid range is **1-60**.
TTL Value	Determines the Time-To-Live /TTL) field value in the IPv4 header. The default value is 64. The valid range is **1-255**.
VID for Source Interface	This field can be used to force the test to use a specific local VLAN interface as the source interface. Leave this field empty for automatic selection based on routing configuration. *Note: You may only specify either the VID or the IP Address for the source interface.*
Source Port Number	This field can be used to force the test to use a specific local interface with the specified port number as the source interface. The specified port must be configured with a suitable IP address. Leave this field empty for automatic selection based on routing configuration. *Note: You may only specify either the Source Port Number or the IP Address for the source interface.*
Address for Source Interface	This field can be used to force the test to use a specific local interface with the specified IP address as the source interface. The specified IP address must be configured on a local interface. Leave this field empty for automatic selection based on routing configuration. *Note: You may only specify either the VID or the IP Address for the source interface.*
Quiet (only print result)	Checking this option will not print the result of each ping request but will only show the final result. After you press **Start**, ICMP packets are transmitted, and the sequence number and round-trip time are displayed upon reception of a reply. The amount of data received inside of an IP packet of type ICMP ECHO_REPLY will always be 8 bytes more than the requested payload data size (the difference is the ICMP header). The page refreshes automatically until responses to all packets are received, or until a timeout occurs. PING 2001::01 (2001::1) from 2001::3: 56 data bytes 64 bytes from 2001::1: seq=0 ttl=64 time=2.118 ms 64 bytes from 2001::1: seq=1 ttl=64 time=2.009 ms 64 bytes from 2001::1: seq=2 ttl=64 time=1.852 ms

	64 bytes from 2001::1: seq=3 ttl=64 time=2.869 ms 64 bytes from 2001::1: seq=4 ttl=64 time=1.845 ms --- 2001::01 ping statistics --- 5 packets transmitted, 5 packets received, 0% packet loss round-trip min/avg/max = 1.845/2.138/2.869 ms

6.3 Traceroute (IPv4)

This page allows you to perform a **traceroute** test over IPv4 towards a remote host. **Traceroute** is a diagnostic tool for displaying the route and measuring transit delays of packets across an IPv4 network.

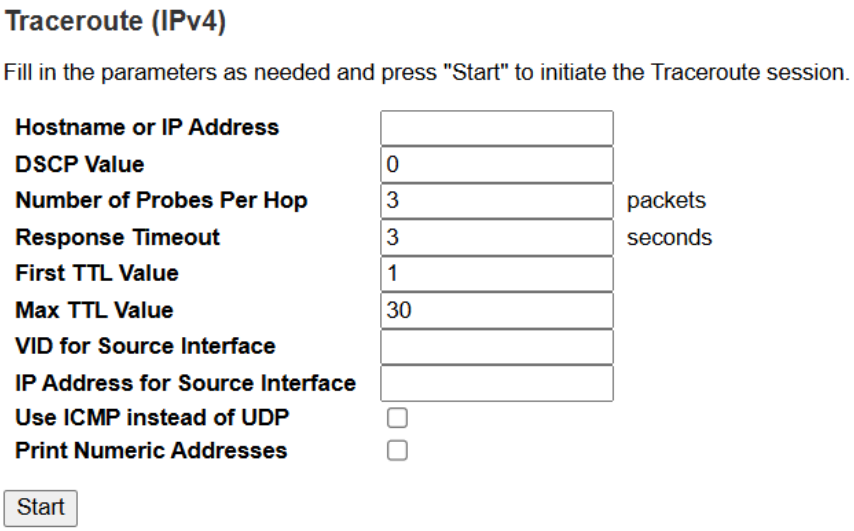

Parameter	Description
Hostname or IP Address	The destination IP Address.
DSCP Value	This value is used for the DSCP value in the IPv4 header. The default value is 0. The valid range is **0-63**.
Number of Probes Per Hop	Determines the number of probes (packets) sent for each hop. The default value is 3. The valid range is **1-60**.
Response Timeout	Determines the number of seconds to wait for a reply to a sent request. The default number is 3. The valid range is **1-86400**.
First TTL Value	Determines the value of the Time-To-Live (TTL) field in the IPv4 header in the first packet sent. The default number is 1. The valid range is **1-30**.
Max TTL Value	Determines the maximum value of the Time-To-Live (TTL) field in the IPv4 header. If this value is reached before the specified remote host is reached the test stops. The default number is 30. The valid range is **1-255**.
VID for Source Interface	This field can be used to force the test to use a specific local VLAN interface as the source interface. Leave this field empty for automatic selection based on routing configuration. Note: You may only specify either the VID or the IP Address for the source interface.
Address for Source Interface	This field can be used to force the test to use a specific local interface with the specified IP address as the source interface. The specified IP address must be configured on a local interface. Leave this field empty for automatic selection based on routing configuration. Note: You may only specify either the VID or the IP Address for the source interface.
Use ICMP instead of UDP	By default, the **traceroute** command will use UDP datagrams. Selecting this option forces it to use ICMP ECHO packets instead.
Print Numeric Address	By default, the **traceroute** command will print out hop information using a reverse DNS lookup for the acquired host ip addresses. This may slow down the display if the DNS information is not

	available. Selecting this option will prevent the reverse DNS lookup and force the **traceroute** command to print numeric IP addresses instead.

6.4 Traceroute (IPv6)

This page allows you to perform a **traceroute** test over IPv6 towards a remote host. **Traceroute** is a diagnostic tool for displaying the route and measuring transit delays of packets across an IPv6 network.

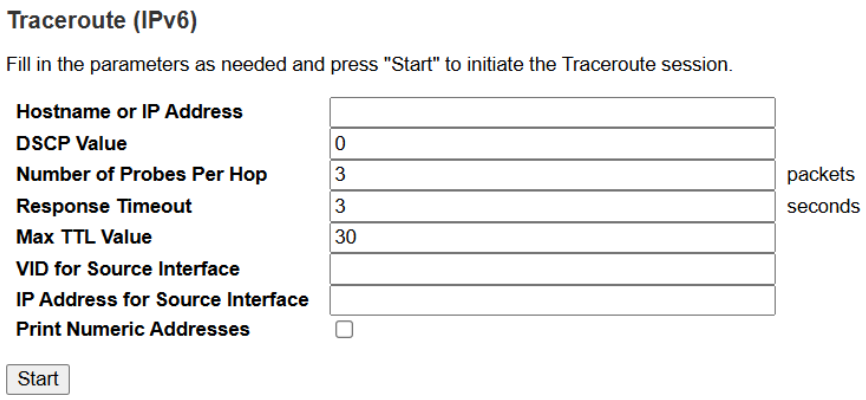

Parameter	Description
Hostname or IP Address	The destination IP Address.
DSCP Value	This value is used for the DSCP value in the IPv4 header. The default value is 0. The valid range is **0-255**.
Number of Probes Per Hop	Determines the number of probes (packets) sent for each hop. The default value is 3. The valid range is **1-60**.
Response Timeout	Determines the number of seconds to wait for a reply to a sent request. The default number is 3. The valid range is **1-86400**.
Max TTL Value	Determines the maximum value of the Time-To-Live (TTL) field in the IPv4 header. If this value is reached before the specified remote host is reached the test stops. The default number is 255. The valid range is **1-255**.
VID for Source Interface	This field can be used to force the test to use a specific local VLAN interface as the source interface. Leave this field empty for automatic selection based on routing configuration. *Note: You may only specify either the VID or the IP Address for the source interface.*
Address for Source Interface	This field can be used to force the test to use a specific local interface with the specified IP address as the source interface. The specified IP address must be configured on a local interface. Leave this field empty for automatic selection based on routing configuration. Note: You may only specify either the VID or the IP Address for the source interface.
Print Numeric Address	By default, the **traceroute** command will print out hop information using a reverse DNS lookup for the acquired host IP addresses. This may slow down the display if the DNS information is not available. Selecting this option will prevent the reverse DNS lookup and force the **traceroute** command to print numeric IP addresses instead.

6.5 Link OAM

6.5.1 MIB Retrieval

This page allows you to retrieve the local or remote OAM MIB variable data on a particular port. Select the appropriate radio button and enter the port number of the switch to retrieve the content of interest.

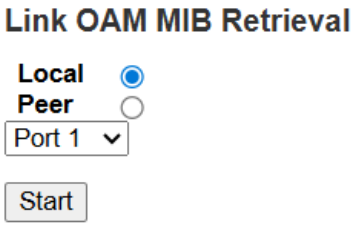

Parameter	Description
Start	Start to retrieve the content.
New Retrieval	New Retrieval to retrieve another content of interest.

6.6 VeriPHY

This page is used for running the VeriPHY Cable Diagnostics for 10/100 and 1G copper ports.

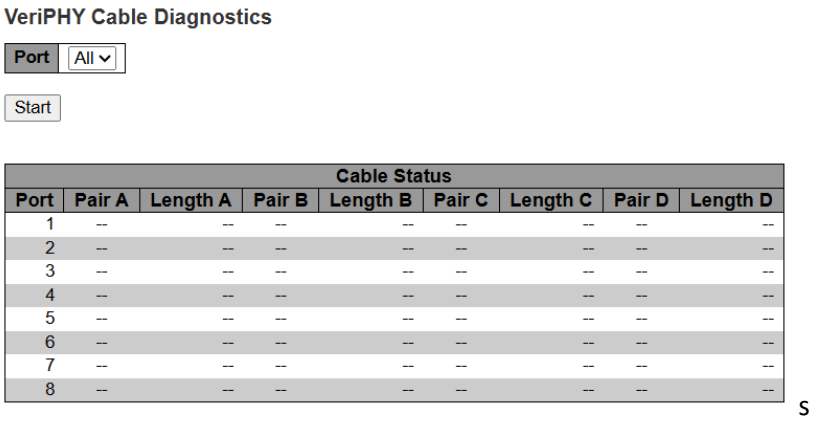

s

Parameter	Description
Start	Press Start to run the diagnostics. This will take approximately 5 seconds. If all ports are selected, this can take approximately 15 seconds. When completed, the page refreshes automatically, and you can view the cable diagnostics results in the cable status table. Note that VeriPHY is only accurate for cables of length 7 - 140 meters. 10 and 100 Mbps ports will be linked down while running VeriPHY. Therefore, running VeriPHY on a 10 or 100 Mbps management port will cause the switch to stop responding until VeriPHY is complete.
Port	The port where you are requesting VeriPHY Cable Diagnostics.
Cable Status	**Port**: Port number. **Pair**: The status of the cable pair. OK - Correctly terminated pair Open - Open pair Short - Shorted pair Short A - Cross-pair short to pair A Short B - Cross-pair short to pair B Short C - Cross-pair short to pair C Short D - Cross-pair short to pair D Cross A - Abnormal cross-pair coupling with pair A Cross B - Abnormal cross-pair coupling with pair B Cross C - Abnormal cross-pair coupling with pair C

	Cross D - Abnormal cross-pair coupling with pair D **Length**: The length (in meters) of the cable pair. The resolution is 3 meters

7 Maintenance

7.1 Restart Device

You can restart the switch on this page. After restart, the switch will boot normally.

Parameter	Description
Yes	Click to restart device.
No	Click to return to the Port State page without restarting.

7.2 Factory Defaults

You can reset the configuration of the switch on this page. Only the IP configuration is retained.
The new configuration is available immediately, which means that no restart is necessary.

Note: Restoring factory default can also be performed by making a physical loopback between port 1 and port 2 within the first minute after the switch reboot. In the first minute after boot, 'loopback' packets will be transmitted at port 1. If a 'loopback' packet is received at port 2, the switch will do a restore to default

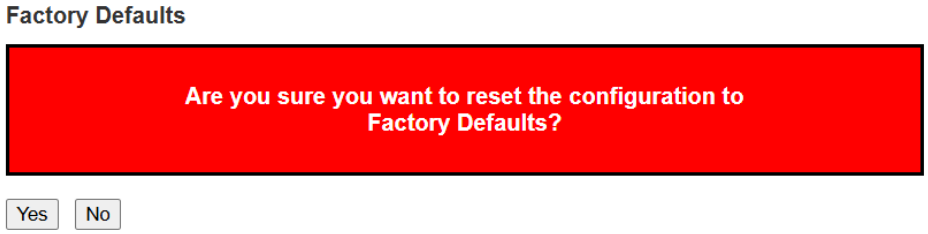

Parameter	Description
Yes	Click to reset the configuration to Factory Defaults.
No	Click to return to the Port State page without resetting the configuration.

7.3 Software

7.3.1 Upload

This page facilitates an update of the firmware controlling the switch.

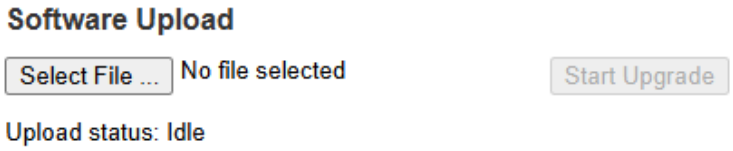

After the software image is uploaded, a page announces that the firmware update is initiated. After about a minute, the firmware is updated and the switch restarts.

Warning*:* While the firmware is being updated, Web access appears to be defunct. The front LED flashes Green/Off with a frequency of 10 Hz while the firmware update is in progress. **Do not restart or power off the device at this time** or the switch may fail to function afterwards.

Parameter	Description
Select File	Browse to the location of a software image and click Start Upgrade.
Start Upgrade	Start the firmware upgrade procedure.

7.3.2 Image Select

This page provides information about the active and alternate (backup) firmware images in the device, and allows you to revert to the alternate image. The web page displays two tables with information about the active and alternate firmware images.

Note:

1. In case the active firmware image is the alternate image, only the "Active Image" table is shown. In this case, the **Activate Alternate Image** button is also disabled.
2. If the alternate image is active (due to a corruption of the primary image or by manual intervention), uploading a new firmware image to the device will automatically use the primary image slot and activate this.
3. The firmware version and date information may be empty for older firmware releases. This does not constitute an error.

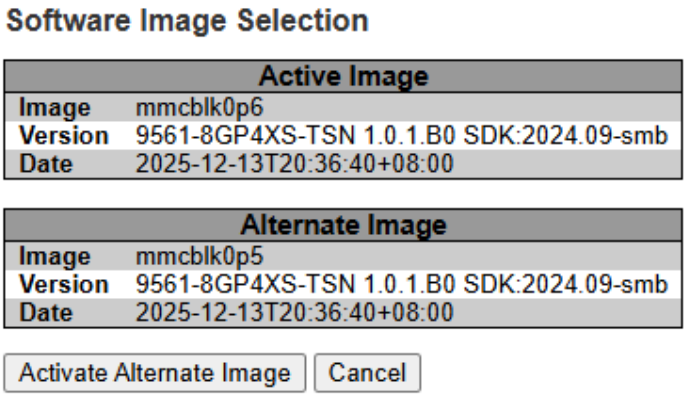

Parameter	Description
Image	The file name of the firmware image, from when the image was last updated.
Version	The version of the firmware image.

Date	The date where the firmware was produced.
Activate Alternate Image	Click to use the alternate image. This button may be disabled depending on system state.
Cancel	Cancel activating the backup image. Navigates away from this page.

7.4 Configuration

7.4.1 Save startup-config

This copies running-config to startup-config, thereby ensuring that the currently active configuration will be used at the next reboot.

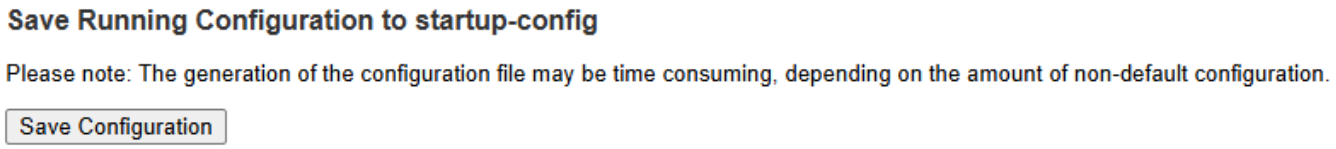

Parameter	Description
Save Configuration	Copy running-config to startup-config

7.4.2 Download

The switch stores its configuration in a number of text files in CLI format. The files are either virtual (RAM-based) or stored in flash on the switch.

The available files are:

- **running-config**: A virtual file that represents the currently active configuration on the switch. This file is volatile.
- **startup-config**: The startup configuration for the switch, read at boot time. If this file doesn't exist at boot time, the switch will start up in default configuration.
- **default-config**: A read-only file with vendor-specific configuration. This file is read when the system is restored to default settings.
- Up to **31** other files, typically used for configuration backups or alternative configurations.

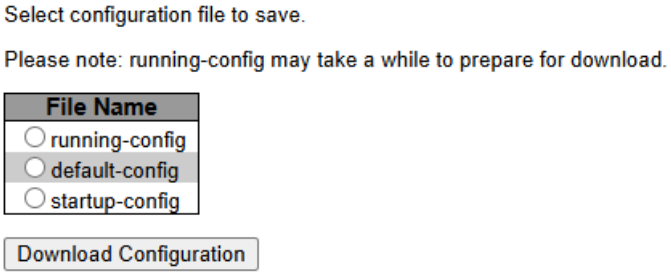

Parameter	Description

Running-config	Select running-config as the file to be downloaded.
Default-config	Select default-config as the file to be downloaded.
Startup-config	Select startup-config as the file to be downloaded.
Download Configuration	Start downloading based on the selected file name.

7.4.3 Upload

It is possible to upload a file from the web browser to all the files on the switch, except the default-config, which is read-only.

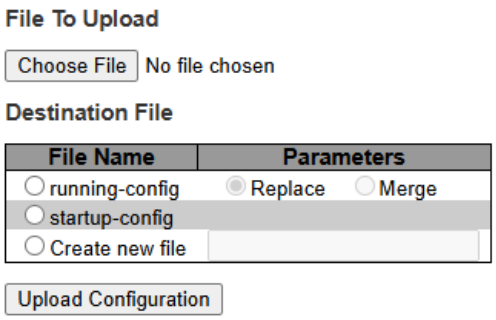

Parameter	Description
Choose File	Browse to the location of a configuration file.
running-config	The destination is running-config, the file will be applied to the switch configuration. This can be done in two ways: **`Replace mode`**: The current configuration is fully replaced with the configuration in the uploaded file. **`Merge mode`**: The uploaded file is merged into *running-config*.
Startup-config	The destination is startup-config; the file will be applied to the switch startup configuration
Create new file	The destination is to create a new file; this file will be uploaded and stored on the switch.
Upload Configuration	Start uploading the selected files to the specified destination.

7.4.4 Activate

It is possible to activate any of the configuration files present on the switch, except for the running-config, which represents the currently active configuration.

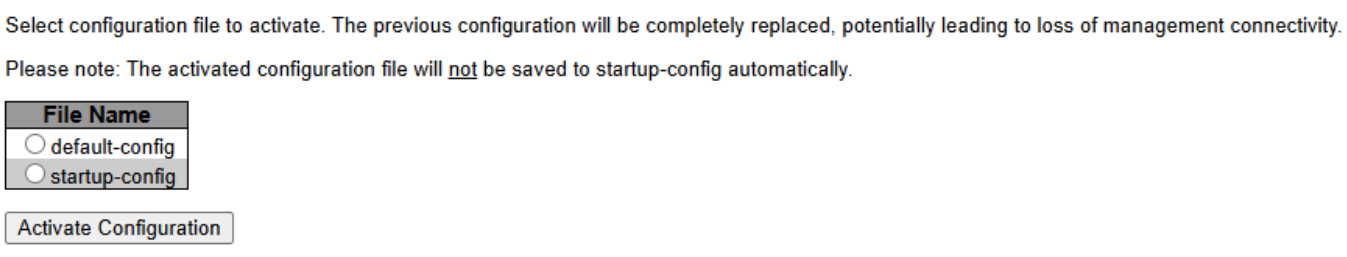

Parameter	Description
default-config	Select the default-config to activate
Startup-config	Select the startup-config to activate
Activate Configuration	It will initiate the process of completely replacing the existing configuration with that of the selected file.

7.4.5 Delete

It is possible to delete any of the writable files stored in flash, including startup-config. If this is done and the switch is rebooted without a prior Save operation, this effectively resets the switch to default configuration.

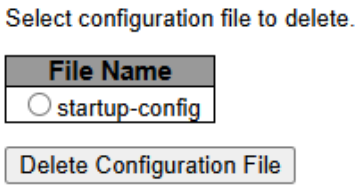

Parameter	Description
Startup-config	Select the startup-config to delete
Delete Configuration file	Start deleting the selected files.

8 Volktek Support

Contact our support services if you have any questions about the switch features or require help with configuration.

Volktek Website: https://www.volktek.com/
FAQ: https://www.volktek.com/?view=faq

8.1 Support

The Support page provides access to Volktek technical support, support contact information, and related online resources.
Scan the QR code or click the QR image to open the Volktek technical support page. An Internet connection is required to open the external support page.
Direct link: Volktek Support (https://www.volktek.com/?view=support).
E-mail: support@volktek.com.
Links: FAQs (https://www.volktek.com/?view=faq) and Volktek Website (https://www.volktek.com/).
Figure: Support page showing QR code, direct support link, e-mail contact, FAQ link, and Volktek website link.

8.2 User Manual

The User Manual page provides a direct product webpage link where users can search for the switch model name and download the user manual.

Click Product webpage to open the Volktek product page (https://www.volktek.com/product_en.php). An Internet connection is required to open the external webpage.

Figure: User Manual page showing the product webpage direct link.