

Liberator 44xx

Web User Guide

LIB-4400 Carrier Ethernet NID



and

LIB-4424 Aggregation Switch



Trademarks

All trademarks and registered trademarks are the property of their respective owners.

Copyright Notice/Restrictions

Copyright © 2013, 2014, 2015, 2016 Net2Edge Limited. All rights reserved.

No part of this work may be reproduced or used in any form or by any means (graphic, electronic or mechanical) without written permission from Net2Edge Limited.

The information contained herein is confidential property of Net2Edge Limited. The use, copying, transfer or disclosure of such information is prohibited except by express written agreement with Net2Edge Limited.

Contact Information

Net2Edge Limited.
Kulite House,
Stroudley Road,
Basingstoke
RG24 8UG
United Kingdom

Tel: +44 345 0130030

Revision History

Rev	Date	Description
A	08/23/2013	Initial release for LIB-4400 Software v 1.6.
B	11/22/2013	Revised for LIB-4400 Software v 1.7 and add LIB-4424
C	01/02/2014	Add DoC (Declaration of Conformity).
D	07/17/2014	Revised for software v 1.9. Adds ZTP Auto Discovery Mode, Performance Monitoring, and Link Fail-over.
E	28/11/2016	Re-Branded manual

Cautions and Warnings

Definitions

Cautions indicate that there is the possibility of poor equipment performance or potential damage to the equipment. Warnings indicate that there is the possibility of injury to a person.

Cautions and Warnings appear here and may appear throughout this manual where appropriate. Failure to read and understand the information identified by this symbol could result in poor equipment performance, damage to the equipment, or injury to persons.

Cautions



Do not ship or store devices near strong electrostatic, electromagnetic, magnetic, or radioactive fields.



Caution: When handling chassis Network Interface Devices (NIDs) observe electrostatic discharge precautions. This requires proper grounding (i.e., wear a wrist strap).



Caution: Copper based media ports, e.g., Twisted Pair (TP) Ethernet, USB, RS232, RS422, RS485, DS1, DS3, Video Coax, etc., are intended to be connected to intra-building (*inside plant*) link segments that are not subject to lightening transients or power faults. They are **not** to be connected to inter-building (*outside plant*) link segments that are subject to lightening.



Caution: **Do not** install the NIDs in areas where strong electromagnetic fields (EMF) exist. Failure to observe this caution could result in poor NID performance.



Caution: Read the installation instructions before connecting the chassis to a power source. Failure to observe this caution could result in poor performance or damage to the equipment.



Caution: Only trained and qualified personnel should install or perform maintenance on the LIB-4400. Failure to observe this caution could result in poor performance or damage to the equipment.



Caution: Do not let optical fibres come into physical contact with any bare part of the body since they are fragile, and difficult to detect and remove from the body.



Caution: Do not bend any part of an optical fibre/cable to a diameter that is smaller than the minimum permitted according to the manufacturer's specification (usually about 65 mm or 2.5 in)!

Warnings



Warning: Use of controls, adjustments or the performance of procedures other than those specified herein may result in hazardous radiation exposure.



Warning: Visible and invisible laser radiation when open. **Do not** look into the beam or view the beam directly with optical instruments. Failure to observe this warning could result in an eye injury or blindness.



Warning: DO NOT connect the power supply module to external power before installing it into the chassis. Failure to observe this warning could result in an electrical shock or death.



Warning: Select mounting bracket locations on the chassis that will keep the chassis balanced when mounted in the rack. Failure to observe this warning could allow the chassis to fall, resulting in equipment damage and/or possible injury to persons.



Warning: Do not work on the chassis, connect, or disconnect cables during a storm with lightning. Failure to observe this warning could result in an electrical shock or death.



Warning: Shock hazard exists if power supply is ejected while powered on.

See "Electrical Safety Warnings" on page 517 for Electrical Safety Warnings translated into multiple languages.

Table of Contents

Trademarks	2
Copyright Notice/Restrictions	2
Contact Information	2
Revision History	2
Cautions and Warnings.....	3
Definitions.....	3
Cautions.....	3
See "Electrical Safety Warnings" on page 530 for Electrical Safety Warnings translated into multiple languages. Table of Contents	4
1. Introduction	8
Document Overview	8
2. Web Interface Menu System	9
Configuration Main Menu.....	10
IP Configuration	12
Detailed Port Statistics.....	317
Monitor > Link OAM > Statistics	325
3. Messages and Troubleshooting.....	427
LIB-44xx Troubleshooting.....	427
Ethernet SAT (Service Activation Testing).....	427
Sync-E and PTP Troubleshooting	428
Sync-E Troubleshooting	428
IEEE 1588 (PTP) Troubleshooting.....	429
EPS Troubleshooting.....	429
ERPS Troubleshooting	430
LIB-44xx Error Recovery	435
Web Interface Messages	437
System Log Messages	488

LIB-44xx Applications Support..... 511

Service..... 512

Warranty 512

Electrical Safety Warnings..... 517

Supported MIBs..... 518

Private MIB OID Assignments 522

Private MIBs 524

A..... 530

B 534

C 534

D..... 538

E 539

F 543

G..... 543

H..... 544

I 544

J 549

L 549

M 551

N..... 553

O..... 555

P 555

Q..... 558

R..... 558

S 560

T 564

U..... 565

V 566

W 568

SAT (Service Activation Testing) Terms 569

10 Gigabit Ethernet Terms 572

Sync-E Terms 576

Figures

Figure 1. SNMP v3 Users, Groups, and Views 45

Figure 2. Spanning Tree Example 119

Figure 3. Multiple Spanning Tree Example 120

Figure 4. ERPS Example 182

Figure 5. Interconnected Ethernet Rings via an Interconnection Node 183

Figure 6. Interconnected Ethernet Rings via Dual Nodes with a Ring Link 183

Figure 7. Ethernet Multi-Ring / Ladder Network 183

Figure 8. Location of RPL for a Sub-ring 184

Figure 9. Intermediate Nodes between Interconnection Nodes 184

Figure 10. Multiple Sub-rings connected to a Major Ring 184

Figure 11. 802.1Q EtherTypes (excerpt from IEEE 802.1ad 2005) 208

Figure 12. All to one bundling VLAN Cases 208

Figure 13. Port Isolation 211

Figure 14a. EPL service per MEF 6.1 220

Figure 14b. EVPL Service per MEF 6.1 220

Figure 15. Bandwidth Profile per UNI 221

Figure 16. Bandwidth Profile per EVC 222

Figure 17. Bandwidth Profile per CoS ID per EVC 222

Figure 18. Provider Bridge E-LINE Service 225

Figure 19. Provider Bridge E-LINE Service 226

Figure 20. Color Aware Dual leaky bucket for Bandwidth Profiling 249

Figure 21. Example SLA for Bandwidth profiling 249

Figure 22. Egress Shaper on Port 1 255

Tables

Table 1: SOAM Terms - IEEE 802.1ag vs. ITU-T 1731 vs. MEF 158

Table 2: SOAM Test Terms - IEEE 802.1ag vs. ITU-T 1731 181

Table 3. EAPOL Counters 341

Table 4. Backend Server Counters 342

Table 5. Last Supplicant/Client Information 343

Table 6: Syslog Info Messages 489

Table 7: Syslog Warning Messages 489

Table 8: Syslog Error Messages 489

Table 9: Public MIBs 518

1. Introduction

Net2Edge's Carrier Ethernet solution delivers the promise of simplicity deployed. This comprehensive solution includes CE 2.0 compliant demarcation devices, access switches, and a service management platform.

The **LIB-4400** is a 10GE Carrier Ethernet NID. The LIB-4400 provides four 1Gbps/10GE SFP+ ports and it includes IEEE 1588v2, SyncE, and Service Activation Test generation. The LIB-4400 CE NID has four 10G SFP+ interfaces with DMI support.

The **LIB-4424** access switch has twenty-four 100/1000Mbps ports and four 10GE uplinks. The LIB-4424 includes IEEE 1588v2, SyncE, and Service Activation Test generation.

The LIB-4400//LIB-4424 software provides a rich set of Carrier Ethernet services, Ethernet switching, and Ethernet transport features. Advanced TCAM-based QoS processing enables delivery of differentiated services with per-service SLA guarantees. Security is assured via separate processing using the LIB-44xx internal processor. The LIB-44xx is designed to support a wide range of MEF-based Carrier Ethernet services for Mobile Backhaul, Business Ethernet, Cloud Assurance and Carrier Exchange E-Access Services.

1.1 Document Overview

The purpose of this manual is to provide information on the configuration, monitoring, diagnostics, and maintenances of the LIB-4400, LIB-4424

The procedures in this manual require you to have completed the install procedure in the LIB-4400 Install Guide manual. This manual documents all of the LIB-44xx models, and notes differences where they apply.

Context-sensitive Help screens are built into the Web interface. A substantial set of technical documents, white papers, case studies, application notes, etc. are available on the Net2Edge Web site at www.Net2Edge.com.

Note that this manual may provide links to third party web sites for which Net2Edge is not responsible.

2. Web Interface Menu System

The LIB-44xx Web interface menu system is shown below in terms of its sub-menus and functions.

Main Menu	Configuration sub-menu	Monitor sub-menu	Diagnostics sub-menu	Maintenance sub-menu
<ul style="list-style-type: none"> ▶ Configuration ▶ Monitor ▶ Diagnostics ▶ Maintenance 	<ul style="list-style-type: none"> ▼ Configuration <ul style="list-style-type: none"> ▶ System ▶ Green Ethernet <ul style="list-style-type: none"> ▪ Thermal Protection ▪ Ports ▶ DHCP ▶ Security ▶ Aggregation ▶ Link OAM <ul style="list-style-type: none"> ▪ Loop Protection ▶ Spanning Tree ▶ IPMC Profile <ul style="list-style-type: none"> ▪ MVR ▶ IPMC ▶ LLDP <ul style="list-style-type: none"> ▪ SyncE ▪ EPS ▪ MEP ▪ ERPS ▪ MAC Table ▶ VLANs ▶ VLAN Translatic ▶ Private VLANs ▶ VCL ▶ Voice VLAN ▶ Ethernet Service ▶ Performance Monitor <ul style="list-style-type: none"> ▶ QoS ▶ HQoS <ul style="list-style-type: none"> ▪ Mirroring ▪ UPnP ▪ PTP ▶ MRP ▶ GVRP <ul style="list-style-type: none"> ▪ sFlow ▶ Traffic Test ▪ UDLD 	<ul style="list-style-type: none"> ▼ Monitor <ul style="list-style-type: none"> ▶ System ▶ Green Ethernet <ul style="list-style-type: none"> ▪ Thermal Protection ▶ Ports ▶ Link OAM <ul style="list-style-type: none"> ▪ Loop Protection ▶ DHCP <ul style="list-style-type: none"> ▪ Security ▶ Aggregation <ul style="list-style-type: none"> ▪ Loop Protection ▶ Spanning Tree <ul style="list-style-type: none"> ▪ MVR ▶ IPMC <ul style="list-style-type: none"> ▪ LLDP ▶ Ethernet Service ▶ Performance Monitor <ul style="list-style-type: none"> ▪ PTP ▪ MAC Table ▶ VLANs <ul style="list-style-type: none"> ▪ MVRP ▪ sFlow ▪ UDLD 	<ul style="list-style-type: none"> ▼ Diagnostics <ul style="list-style-type: none"> ▪ Ping ▶ Link OAM <ul style="list-style-type: none"> ▪ Ping6 ▪ VeriPHY 	<ul style="list-style-type: none"> ▼ Maintenance <ul style="list-style-type: none"> ▪ Restart Device ▪ Factory Default ▼ Software <ul style="list-style-type: none"> ▪ Upload ▪ Image Select ▼ Configuration <ul style="list-style-type: none"> ▪ Save startup-c ▪ Download ▪ Upload ▪ Activate ▪ Delete

The four Main Menu selections are:

Configuration - lets you define system operating parameters for the available LIB-44xx features.

Monitor - lets you view and track the LIB-44xx operating functions. See ‘Monitor’ on page 310.

Diagnostics - provides access to the full set of LIB-44xx tests and verification functions. See ‘Diagnostics Main Menu’ on page 411.

Maintenance - supports the LIB-44xx troubleshooting and service functions. See ‘Maintenance Menu’ on page 417.

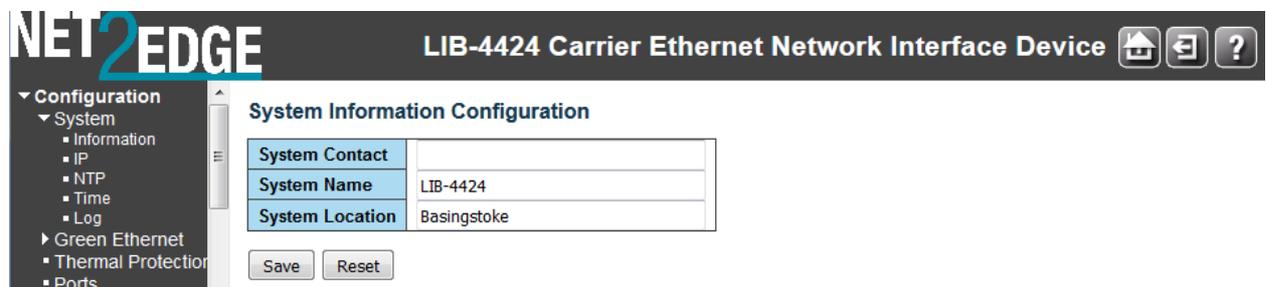
Click one or more of the main menu selections to display its sub-menus.

Each of these sub-menus and their functions are described in the following sections.

2.1 Configuration Main Menu

Configuration > System

The LIB-44xx system information is configured from the **Configuration > System** menu path. Here you can configure LIB-44xx device level Information, IP, NTP, time, logging, and Zero Touch Provisioning.



The LIB-44xx system information parameters are explained below:

System Contact

Enter the textual identification of the contact person for this managed node, together with information on how to contact this person. The valid string length is 0 to 255 characters, and the allowed content is the set of ASCII characters from 32 to 126. In the ASCII code table, the characters from 32 to 126 inclusive are printable. (The ASCII characters from 0 to 32 and 127 are defined as control characters and are not printable.) This field is blank by default.

If you delete an existing System Contact entry, the message “*System Contact is empty. Do you want to proceed anyway?*” displays. Verify the action and continue operation.

System Name

Enter 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-Z, a-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 valid string length is 0 to 255 characters. This field is blank by default.

If you delete an existing System Name entry, the message “*System Name is empty. Do you want to proceed anyway?*” displays. Verify the action and continue operation.

System Location

Enter the physical location (wiring closet, floor, or building) of this node (e.g., telephone closet, 3rd floor). The valid string length is 0 to 255 characters, and the allowed content is the set of ASCII characters from 32 to 126. In the ASCII code table, the characters from 32 to 126 inclusive are printable. (The ASCII characters from 0 to 32 and 127 are defined as control characters and are not printable.) This field is blank by default.

If you delete an existing System Location entry, the message “*System Location is empty. Do you want to proceed anyway?*” displays. Verify the action and continue operation.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

An updated System Information Configuration table is shown below:

NET2EDGE LIB-4424 Carrier Ethernet Network Interface Device

Configuration

- System
 - Information
 - IP
 - NTP
 - Time
 - Log
- Green Ethernet
- Thermal Protection
- Ports

System Information Configuration

System Contact	Site1
System Name	LIB-4424
System Location	Basingstoke

Save Reset

Note that changes in the Web interface here will change the CLI prompt appearance. Based on the screen above, the CLI prompt will change to include the newly added System Name:

from: >
to: SCENID010 : />

If you enter a space character or other invalid entry and click the **Save** button, the message “*System Name' is not valid. Please refer to the help page for the valid format.*” displays.

NET2EDGE LIB-4424 Carrier Ethernet Network Interface Device

System Information Configuration

'System Name' is not valid. Please refer to the help page for the valid format.

OK

Click the **OK** button to clear the webpage message, enter the System Name without any spaces, and click the **Save** button.

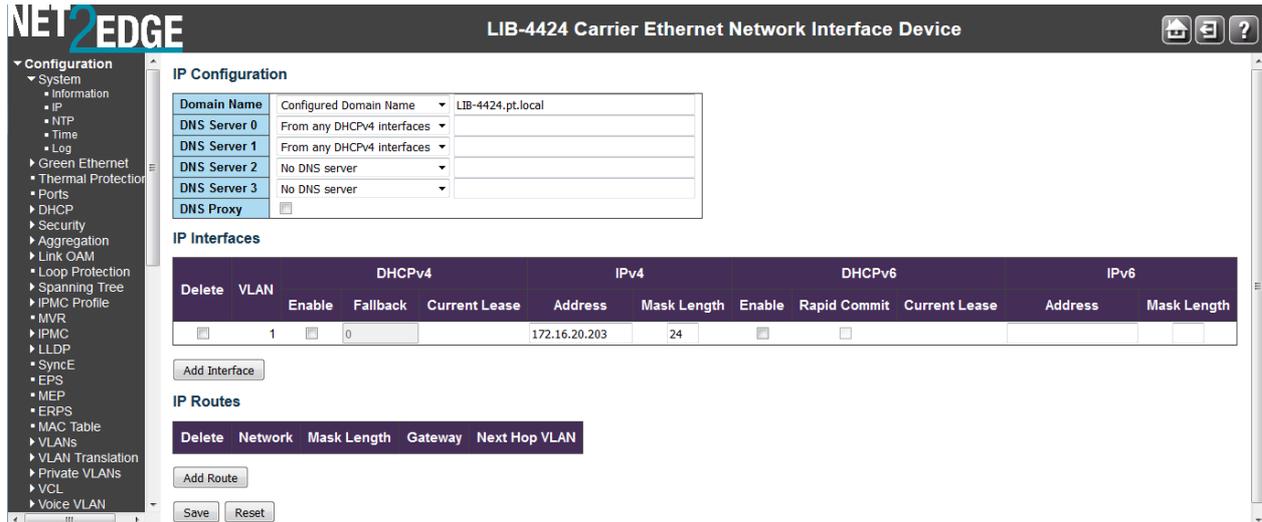
Character Support Note

The LIB-44xx supports the Space * < > : % / \ " ? |, and all other keyboard characters. The LIB-4424 supports an individual "dot" character.

IPv4 Configuration

Configure the LIB-44xx-managed IPv4 information from **Configuration > System > IP** as one of the first steps.

The LIB-44xx supports an IPv4 / IPv6 dual stack. The LIB-44xx can be assigned IP address statically or dynamically using DHCP.



The **Configured** column is used to view or change the IP configuration. The **Current** column displays the active IP configuration.

2.2 IP Configuration

DHCP Client

Enable the DHCP client by checking this box. If DHCP fails and the configured IP address is zero, DHCP will retry. If DHCP fails and the configured IP address is non-zero, DHCP will stop and the configured IP settings will be used. The DHCP client will announce the configured System Name as hostname to provide DNS lookup.

DHCP: some of the things to consider when DHCP is invoked:

1. When DHCP is enabled, the device fails to contact the DHCP server after retries; the last configured static IP address will take effect.
2. If the DHCP lease time expires, the dynamic address assigned should not be used and it should default to last known static IP address.
3. Configuration backup should not backup the leased IP address; only the DHCP state should be backed up.
4. The LIB-44xx uses the MAC Address in DHCP option 61 client-identifier. IETF RFC 2132 defines DHCP Options and BOOTP Vendor Extensions at <http://www.ietf.org/rfc/rfc2132.txt>. Per RFC section 9.14. Client-identifier: "This option is used by DHCP clients to specify their unique identifier. DHCP servers use this value to index their database of address bindings. This value is expected to be unique for all clients in an administrative domain. Vendors and

system administrators are responsible for choosing client-identifiers that meet this requirement for uniqueness. The code for this option is 61, and its minimum length is 2."

DNS: The LIB-44xx supports Domain Name Service and it can be given a hostname instead of IP address.

A DNS server address for name resolution must be provided. The LIB-44xx can act in a proxy role for other nodes that are connected to it by passing on the DNS request/response to and from the server. If you check the DHCP Client Configured checkbox on this page, you may lose IP connectivity. This could in turn require reconfiguration of this PC. Make sure this is what you want before you **Save** to continue.

You can click the **Renew** button to renew the IPv4 configuration. This button is only available if the DHCP Client Configured checkbox is checked, otherwise it is greyed out.

IP Address

Provide the IPv4 address of this LIB-44xx in dotted decimal notation (e.g., 192.168.1.110).

IP Mask

Provide the IPv4 mask of this LIB-44xx in dotted decimal notation (e.g., 255.255.255.0).

IP Router

Provide the IPv4 address of the router in dotted decimal notation (e.g., 192.168.1.1).

VLAN ID

Provide the managed VLAN ID. A VLAN ID is a 12-bit field specifying the VLAN to which the frame belongs. The valid range is 1 to 4094 for this Management VLAN ID. This Management VLAN provides a secure channel for all management traffic to/from the device.

Note: Be sure to set the Management VLAN configuration before you set the VLAN configuration. The Management VLAN configured from the web at "**Configuration > System > IP: VLAN ID**" should be one of the first steps in the overall LIB-44xx configuration process. The Management VLAN is configured from the CLI with the "IP MVlan" commands.

The Management VLAN is used to establish an IP connection to the switch from a workstation connected to a port in the VLAN. Management VLAN is used for Out-of-band management of the LIB-44xx from a remote location using a protocol such as telnet, SSH, or SNMP.

Normally, the Management VLAN is VLAN 1, but you can configure any VLAN as the Management VLAN. recommends not using VLAN 1, and also not to use any VLAN that carries user data traffic as Management VLAN. You must configure an IP address and a default gateway for Management VLAN.

After configuring IP address and default gateway for Management VLAN, you can telnet or SSH to the switch to perform switch management functions.

By default, all LIB-44xx ports belong to default VLAN 1 and when a management IP address is assigned to the LIB-44xx, it is accessible by all ports belonging to the default VLAN.

(Note that this "Management VLAN" function is not to be confused with the "Management Port" which is configured at the **Configuration > VLANs > Ports: Management Port - PortType** menu path).

DNS Server

Provide the IP address of the DNS Server in dotted decimal notation.

IP DNS Proxy Configuration

When **DNS Proxy** is enabled, the LIB-44xx will relay DNS requests to the currently configured DNS server on the LIB-44xx, and reply as a DNS resolver to the client device on the network.

Note that setting these fields does not provide the full set of IP, BootP, VLAN, DNS server, and Management VLAN / member ports configuration. See the related Configuration sections of this manual for additional information.

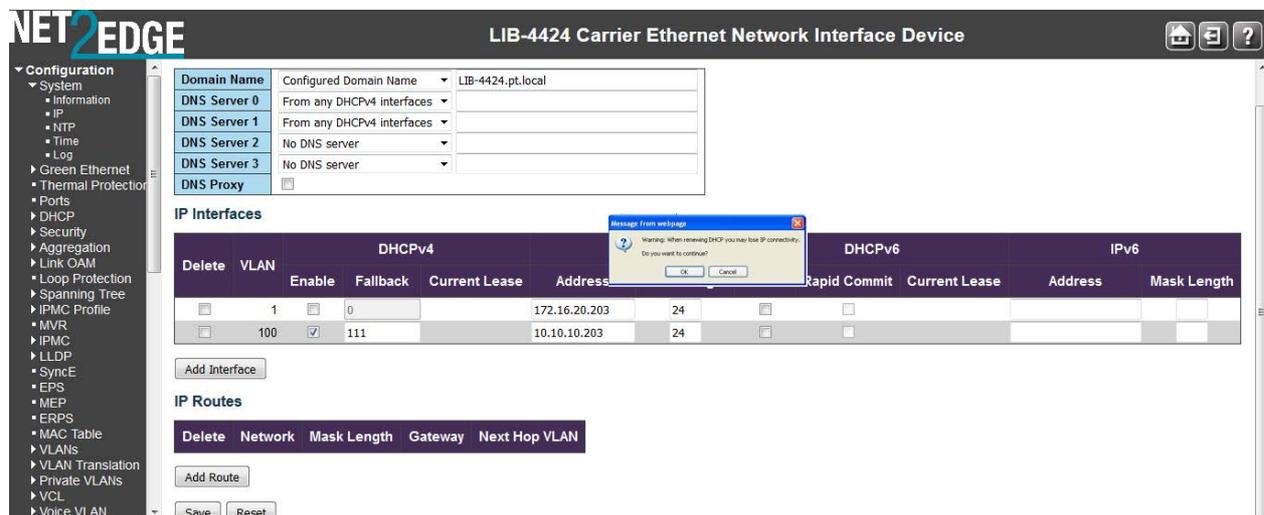
Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Renew: Click to renew DHCP. This button is only available if DHCP Client is enabled ('Configured' checkbox checked) and the configuration is Saved.

After you click the **Renew** button, the message "Warning: When renewing DHCP you may lose IP connectivity. Do you want to continue?" displays.



If you are sure you want to do this, click the **OK** button to clear the webpage message.

If you do not want to renew DHCP / lose the IP connection, click the **Cancel** button and continue operation.

After you click the **Save** button, the message "Warning: When changing parameters on this page, you may lose IP connectivity. This could in turn require reconfiguration of this PC. Do you want to continue?" may display.

The screenshot shows the configuration page for a LIB-4424 Carrier Ethernet Network Interface Device. The left sidebar contains a navigation menu with categories like Configuration, System, Information, IP, NTP, Time, Log, Green Ethernet, Thermal Protection, Ports, DHCP, Security, Aggregation, Link OAM, Loop Protection, Spanning Tree, IPMC Profile, MVR, IPMC, LLDP, SyncE, EPS, MEP, ERPS, MAC Table, VLANs, VLAN Translation, Private VLANs, VCL, and Voice VLAN.

The main content area is divided into several sections:

- Domain Name:** Configured Domain Name is LIB-4424.pt.local.
- DNS Servers:** DNS Server 0, 1, 2, and 3 are configured. DNS Server 0 and 1 are set to 'From any DHCPv4 interfaces'. DNS Server 2 and 3 are set to 'No DNS server'. A checkbox for 'DNS Proxy' is present.
- IP Interfaces:** A table with columns for Delete, VLAN, DHCPv4 (Enable, Fallback, Current Lease), IPv4 (Address, Mask Length), DHCPv6 (Enable, Rapid Commit, Current Lease), and IPv6 (Address, Mask Length). Two interfaces are listed: VLAN 1 and VLAN 100. A warning message is overlaid on this table.
- IP Routes:** A table with columns for Delete, Network, Mask Length, Gateway, and Next Hop VLAN. An 'Add Route' button is visible.

Buttons for 'Add Interface', 'Add Route', 'Save', and 'Reset' are also present.

If you are sure you want to change these parameters, click the **OK** button to clear the webpage message. Follow any on-screen messages.

If you do not want to change parameters, click the **Cancel** button and continue operation.

IPv6 Configuration

Configure the LIB-44xx-managed IPv6 information on this page from Configuration > System > IPv6. The LIB-44xx supports IPv4 / IPv6 dual stack. The LIB-44xx can be assigned IP address statically or dynamically using DHCP.

IPv6 Configuration Considerations

For the latest feature information and caveats, see the release notes for your particular device and soft-ware release. The prerequisites and restrictions below apply to all LIB-44xx models unless otherwise noted.

When changing certain parameters on this page, you may lose IP connectivity. This could in turn require reconfiguration of this PC. Determine whether your particular computer will require reconfiguration. For example, for Microsoft .NET Framework version 2.0 and later, IPv6 is enabled by default. For .NET Framework version 1.1 and earlier, IPv6 is disabled by default. For more information see the MSDN article at <http://msdn.microsoft.com/en-us/library/8db2058t.aspx>.

IETF RFC 2462 defines both a stateful and stateless address autoconfiguration mechanism for IPv6. Stateless autoconfiguration requires no manual configuration of hosts, minimal (if any) configuration of routers, and no additional servers.

Stateless autoconfiguration allows a host to generate its own addresses using a combination of locally available information and information advertised by routers. Routers advertise prefixes that identify the subnet(s) associated with a link, while hosts generate an "interface identifier" that uniquely identifies an interface on a subnet. An address is formed by combining the two. In the absence of routers, a host can only generate link-local addresses. However, link-local addresses are sufficient for allowing communication among nodes attached to the same link. Stateless auto-configuration is suitable for small organizations and individuals. It allows each host to determine its address from the contents of received user advertisements. It makes use of the IEEE EUI-64 standard to define the network ID portion of the address. The stateless approach is used when a site

is not particularly concerned with the exact addresses hosts use, so long as they are unique and properly routable. Stateless auto-configuration is suitable for small organizations and individuals. In this case, each host determines its addresses from the contents of received router advertisements. Using the IEEE EUI-64 standard to define the network ID portion of the address, it is reasonable to assume the uniqueness of the host address on the link.

Stateful autoconfiguration has hosts obtain interface addresses and/or configuration information and parameters from a server. Servers maintain a database that keeps track of which addresses have been assigned to which hosts. The stateful autoconfiguration protocol allows hosts to obtain addresses, other configuration information or both from a server. Stateful auto-configuration requires a certain level of human intervention because it needs a Dynamic Host Configuration Protocol for IPv6 (DHCPv6) server for the installation and administration of the nodes. The DHCPv6 server keeps a list of nodes to which it supplies configuration information. It also maintains state information so the server knows how long each address is in use, and when it might be available for reassignment. The stateful approach is used when a site requires tighter control over exact address assignments.

Both stateful and stateless address autoconfiguration may be used simultaneously. Stateful auto-configuration requires some human intervention as it makes use of the Dynamic Host Configuration Protocol for IPv6 (DHCPv6) for installation and administration of nodes over a network. The DHCPv6 server maintains a list of nodes and the information about their state to know the availability of each IP address from the range specified by the network administrator. The autoconfiguration process specified in RFC 2462 applies only to hosts and not routers. Since host autoconfiguration uses information advertised by routers, routers must be configured by some other means. However, it is expected that routers will generate link-local addresses using the mechanism described in the RFC. Also, routers are expected to successfully pass the Duplicate Address Detection procedure on all addresses prior to assigning them to an interface.

Stateless and stateful autoconfiguration complement each other. For example, a host can use stateless autoconfiguration to configure its own addresses, but use stateful autoconfiguration to obtain other information. The site administrator specifies which type of autoconfiguration to use through the setting of appropriate fields in Router Advertisement messages (Discovery). Stateful autoconfiguration for IPv6 is the subject of DHCPv6.

Configure the LIB-44xx IPv6 information on this page from **Configuration > System > IPv6**.

Note: changes made on this page may cause a disruption in IP connectivity.

The screenshot shows the NET2EDGE web interface for a LIB-4424 Carrier Ethernet Network Interface Device. The main content area is titled "IP Configuration" and contains several sections:

- IP Configuration Form:** Includes fields for "Domain Name" (Configured Domain Name: LIB-4424.pt.local), "DNS Server 0", "DNS Server 1", "DNS Server 2", "DNS Server 3", and "DNS Proxy".
- IP Interfaces Table:** A table with columns for "Delete", "VLAN", "DHCPv4", "IPv4", "DHCPv6", and "IPv6". The "DHCPv4" section includes "Enable", "Fallback", and "Current Lease". The "IPv4" section includes "Address" and "Mask Length". The "DHCPv6" section includes "Enable", "Rapid Commit", and "Current Lease". The "IPv6" section includes "Address" and "Mask Length".
- IP Routes Table:** A table with columns for "Delete", "Network", "Mask Length", "Gateway", and "Next Hop VLAN".

The **Configured** column is used to view or change the IPv6 configuration (read-write).

The **Current** column shows the active IPv6 configuration (read only).

Auto Configuration

Enable IPv6 auto-configuration by checking this box. If this fails, the configured IPv6 address is zero. The router may delay responding to a router solicitation for a few seconds, the total time needed to complete auto-configuration can be significantly longer.

When changing parameters on this page, you may lose IP connectivity. This could in turn require reconfiguration of this device. Make sure this is what you want before you click **Save** to continue.

You can click the **Renew** button to renew the IPv6 Auto Configuration. This button is only available if IPv6 Auto Configuration is enabled, otherwise it is greyed out.

Address

Provide the IPv6 address of this LIB-44xx. An IPv6 address is in 128-bit records represented as eight fields of up to four hexadecimal digits with a colon separating each field (e.g.).

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 only appear once.

It can also represent a legally valid IPv4 address ('**::192.1.2.34**').

In IPv6, a valid address can be a preferred or deprecated address. A valid address may appear as the source or destination address of a packet, and the internet routing system is expected to deliver packets sent to a valid address to their intended recipients.

Prefix

Provide the IPv6 Prefix of this LIB-44xx. The valid range is **1** to **128**.

In IPv6, routers advertise prefixes that identify the subnet(s) associated with a link, while hosts generate an "interface identifier" that uniquely identifies an interface on a subnet. An address is formed by combining the two. In the absence of routers, a host can only generate link-local addresses. However, link-local addresses are sufficient for allowing communication among nodes attached to the same link.

Routers generate periodic Router Advertisements that include options listing the set of active prefixes on a link. A 'Lease lifetime' provides the mechanism through which a site phases out old prefixes.

Router

Provide the IPv6 gateway address of this LIB-44xx. 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 short way to represent multiple 16-bit groups of contiguous zeros; but it can only appear once. It can also represent a valid IPv4 address (e.g., '::192.1.2.34').

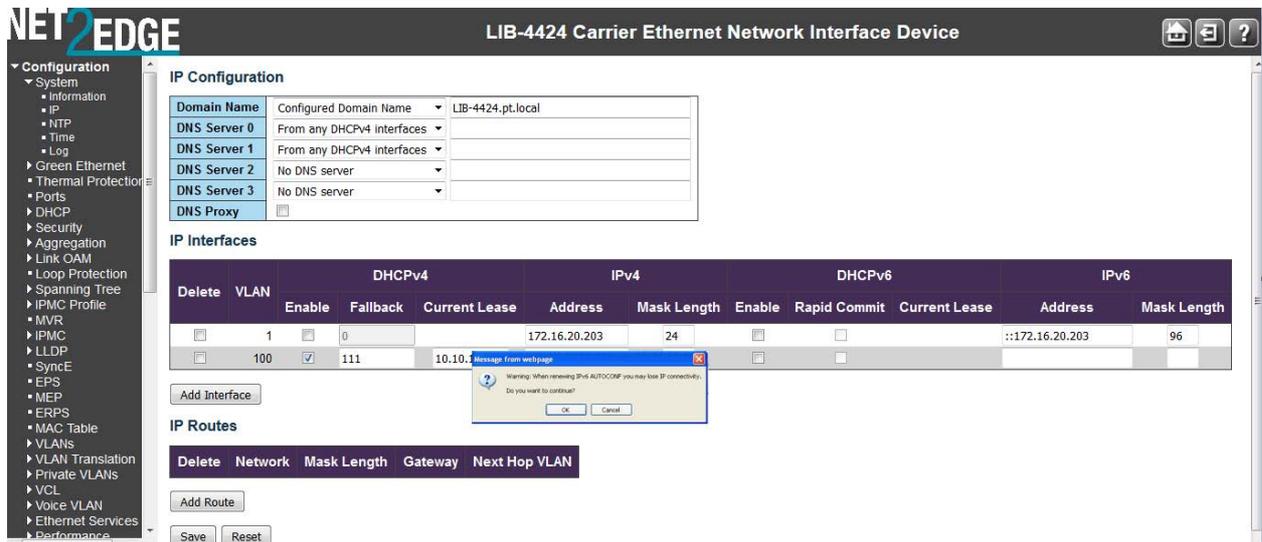
Buttons

Save: Click to save changes.

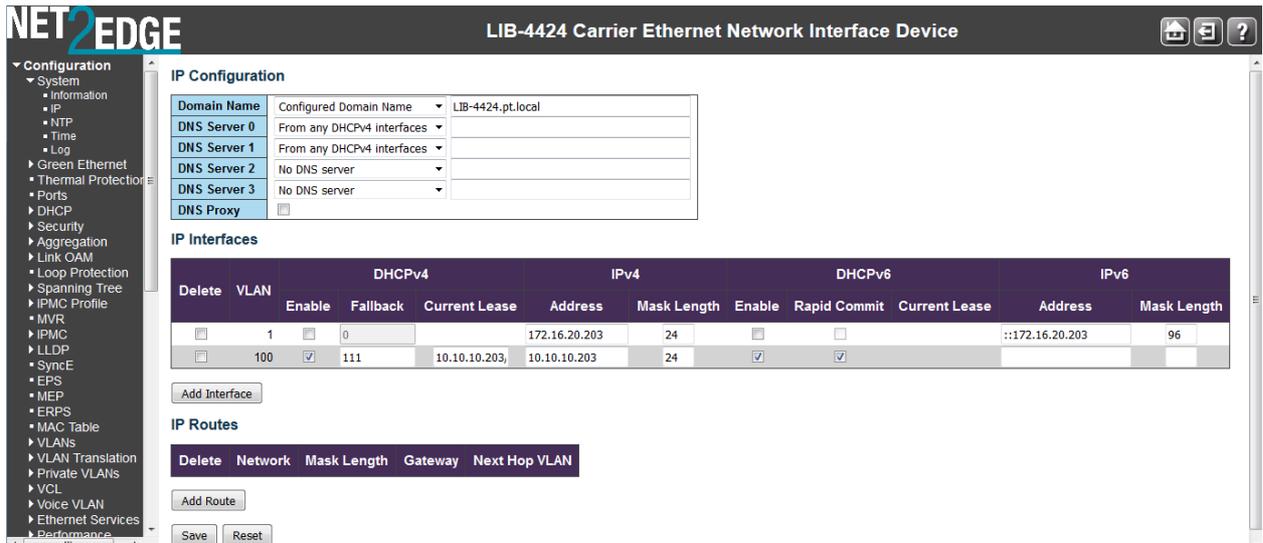
Reset: Click to undo any changes made locally and revert to previously saved values.

Renew: Click to renew the IPv6 Auto Configuration. This button is only available if IPv6 Auto Configuration is enabled, otherwise it is greyed out.

Warning: If you click the **Renew** button, a warning message displays. Click the **'Cancel'** button if you are not sure you want to renew the IPv6 Auto Configuration. Click the **'OK'** button only if you are sure you want to renew the IPv6 Auto Configuration, and understand that the current LIB-44xx web session may drop.

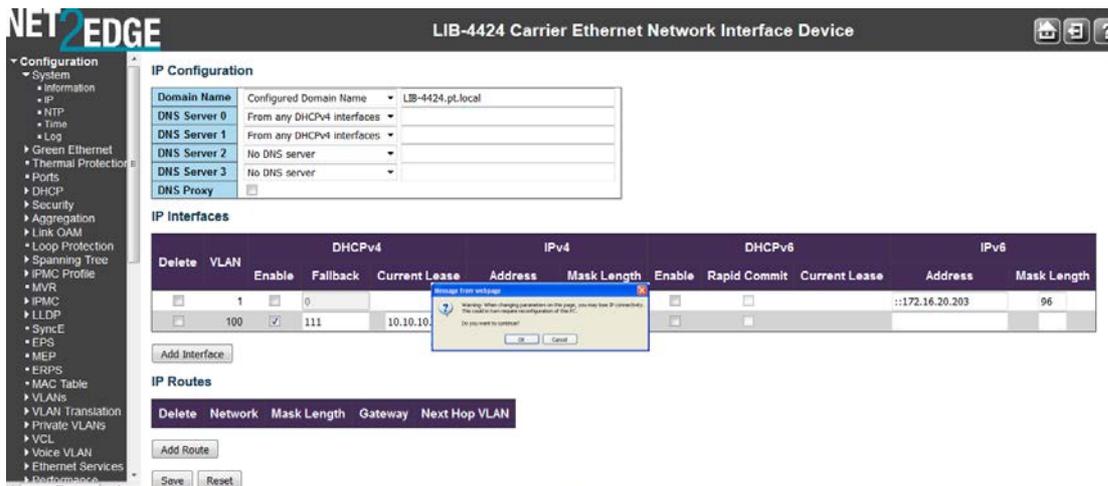


After renewing IPv6 AUTOCONF, the IPv6 Configuration table displays again with renewed information.



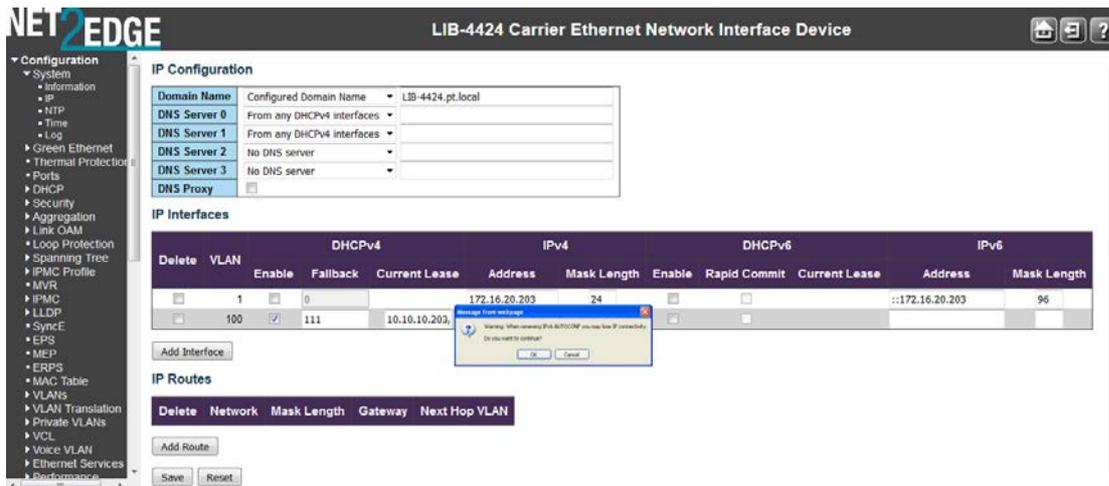
Uncheck the ‘Auto Configuration Configured’ checkbox, click ‘Save’, check the ‘Auto Configuration Configured’ checkbox, and click ‘Save’ to renew again with auto config information.

With the Configured checkbox unchecked, after you click the **Save** button, the message “Warning: When changing parameters on this page, you may lose IP connectivity. This could in turn require reconfiguration of this PC. Do you want to continue?” may display.



If you are sure you want to change these parameters, click the **OK** button to clear the webpage message. Follow any on-screen messages.

If you do not want to change parameters, click the **Cancel** button and continue operation. When you click the **Renew** button, the message “Warning: When renewing IPv6 AUTOCONF you may lose IP connectivity. Do you want to continue?” may display.



If you are sure you want to change these parameters, click the **OK** button to clear the webpage message. Follow any on-screen messages.

If you do not want to change parameters, click the **Cancel** button and continue operation.

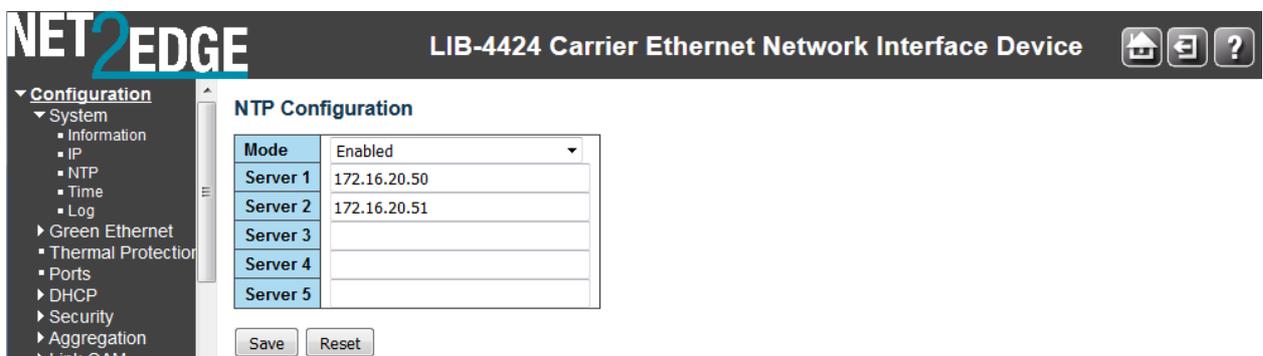
Messages:

*Parameter <server_ipv6> doesn't allowed all zero or all 'ff' Sever already exist! Delete it first.
Using IPv6 multicast address is not allowed here.*

NTP Configuration

Configure NTP on this page from **Configuration > System > NTP**. Network Time Protocol (NTP) is a network protocol for synchronizing the clocks of computer systems. NTP uses UDP (datagrams) as transport layer.

The LIB-44xx uses NTP for real time clock synchronization with the network time server. The NTP is compliant with RFC 5905. The LIB-44xx takes care of day light saving options where used. The management interfaces provide options to configure NTP and report the network time synchronized. The LIB-44xx time obtained is used for all LIB-44xx services that need a timestamp.



Mode

Indicates the NTP mode operation. Possible modes are:
Enabled: Enable NTP mode operation. When NTP mode operation is enabled, the agent forwards

NTP messages between the clients and the server when they are not on the same subnet domain.
Disabled: Disable NTP mode operation.

Server

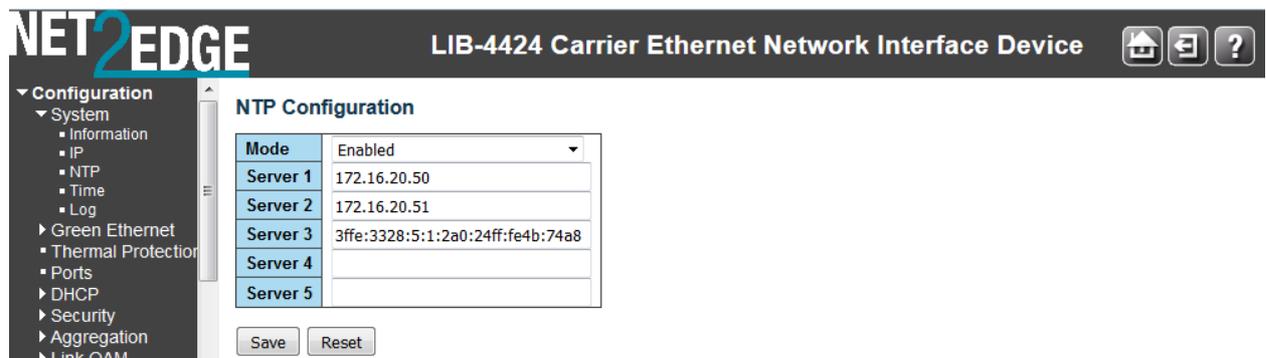
Provide the NTP IPv4 or IPv6 address of this LIB-44xx. IPv6 address is in 128-bit records represented as eight fields of up to four hexadecimal digits with a colon separating each field (:). For example, 'fe80::215:c5ff:fe03:4dc7'. The symbol '::' is a special syntax that can be used as a shorthand way of representing multiple 16-bit groups of contiguous zeros; but it can only appear once. It can also represent a legal, valid IPv4 address. For example, '::192.1.2.34'. Not that you can 'cut and paste' information to and from this field.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

The example below shows NTP enabled with one IPv4 NTP server and one IPv6 NTP server.



Messages:

Parameter <server_ipv6> doesn't allowed all zero or all 'ff'

Sever already exist! Delete it first.

Using IPv6 multicast address is not allowed here.

Problem: NTP does not work with Windows Server 2008

Meaning: Microsoft's W32Time service cannot reliably maintain sync time to the range of 1 to 2 seconds needed for high accuracy environments. The W32Time service is not a full-featured NTP solution that meets time-sensitive application needs.

Recovery: See the Microsoft Support site at <http://support.microsoft.com/kb/939322>.

Time Configuration

page allows you to configure the Time Zone and Daylight Savings Time (DST) parameters from the Configuration > System > Time menu path.

NET2EDGE LIB-4424 Carrier Ethernet Network Interface Device

Time Zone Configuration

Time Zone	None
Acronym	(0 - 16 characters)

Daylight Saving Time Configuration

Daylight Saving Time	Disabled
-----------------------------	----------

Start Time settings

Month	Jan
Date	1
Year	2014
Hours	0
Minutes	0

End Time settings

Month	Jan
Date	1
Year	2097
Hours	0
Minutes	0

Offset settings

Offset	1 (1 - 1440) Minutes
---------------	----------------------

Save Reset

The Time Zone and Daylight Savings Time (DST) parameters are described below:

Time Zone Configuration

Time Zone

Time Zone: Lists various Time Zones worldwide. Select the appropriate Time Zone from the drop down and click Save to set. See the end of this section for the Time Zone selections.

Acronym: User can set the acronym of the time zone. This is a User configurable acronym to identify the time zone. Enter up to 16 alpha-numeric characters including '-', ' ' or '.' characters.

Daylight Saving Time Configuration

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. The default is **Disabled**.

Select '**Disabled**' 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.

Recurring Configurations

Start time settings

Week - Select the starting week number (1-5).

Day - Select the starting day (**Sun**, **Mon**, **Tue**, **Wed**, **Thu**, **Fri**, or **Sat**).

Month - Select the starting month (**Jan**, **Feb**, **Mar**, **Apr**, **May**, **Jun**, **Jul**, **Aug**, **Sep**, **Oct**, **Nov**, or **Dec**).

Hours - Select the starting hour (**0** - **23**).

Minutes - Select the starting minute (**0** - **59**).

End time settings

Week - Select the ending week number (1-5).

Day - Select the ending day (**Sun**, **Mon**, **Tue**, **Wed**, **Thu**, **Fri**, or **Sat**).

Month - Select the ending month (**Jan**, **Feb**, **Mar**, **Apr**, **May**, **Jun**, **Jul**, **Aug**, **Sep**, **Oct**, **Nov**, or **Dec**).

Hours - Select the ending hour (**0** - **23**).

Minutes - Select the ending minute (**0** - **59**).

Offset settings

Offset - Enter the number of minutes to add during Daylight Saving Time. (Range: **1** to **1440**)

Non Recurring Configurations

Start time settings

Month - Select the starting month (**Jan**, **Feb**, **Mar**, **Apr**, **May**, **Jun**, **Jul**, **Aug**, **Sep**, **Oct**, **Nov**, or **Dec**)..

Date - Select the starting date (**1** - **31**).

Year - Select the starting year (**2000** - **2097**).

Hours - Select the starting hour (**0** - **23**).

Minutes - Select the starting minute (**0** - **59**).

End time settings

Month - Select the ending month (**Jan**, **Feb**, **Mar**, **Apr**, **May**, **Jun**, **Jul**, **Aug**, **Sep**, **Oct**, **Nov**, or **Dec**).

Date - Select the ending date (**1** - **31**).

Year - Select the ending year (**2000** - **2097**).

Hours - Select the ending hour (**0** - **23**).

Minutes - Select the ending minute (**0** - **59**).

Offset settings

Offset - Enter the number of minutes to add during Daylight Saving Time (**1** - **1440**).

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Time Zones List

The various Time Zone selections available worldwide are listed below:

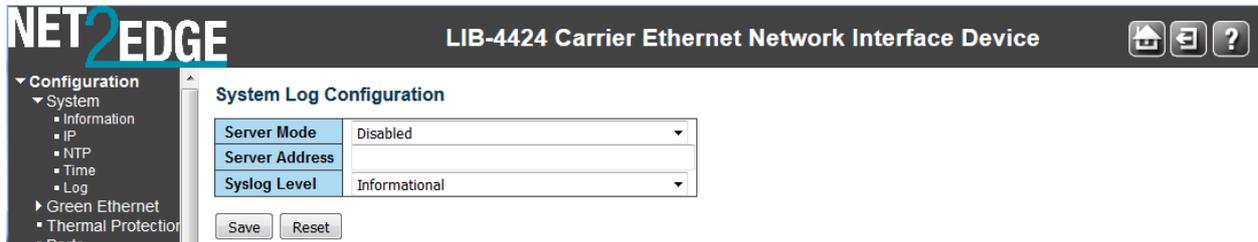
None

(GMT-12:00) International Date Line West
 (GMT-11:00) Midway Island, Samoa
 (GMT-10:00) Hawaii
 (GMT-09:00) Alaska
 (GMT-08:00) Pacific Time (US and Canada)
 (GMT-08:00) Tijuana, Baja California
 (GMT-07:00) Arizona
 (GMT-07:00) Chihuahua, La Paz, Mazatlan - New
 (GMT-07:00) Chihuahua, La Paz, Mazatlan - Old
 (GMT-07:00) Mountain Time (US and Canada)
 (GMT-06:00) Central America
 (GMT-06:00) Central Time (US and Canada)
 (GMT-06:00) Guadalajara, Mexico City, Monterrey - New
 (GMT-06:00) Guadalajara, Mexico City, Monterrey - Old
 (GMT-06:00) Saskatchewan
 (GMT-05:00) Bogota, Lima, Quito, Rio Branco
 (GMT-05:00) Eastern Time (US and Canada)
 (GMT-05:00) Indiana (East)
 (GMT-04:30) Caracas
 (GMT-04:00) Atlantic Time (Canada)
 (GMT-04:00) La Paz
 (GMT-04:00) Manaus
 (GMT-04:00) Santiago
 (GMT-03:30) Newfoundland
 (GMT-03:00) Brasilia
 (GMT-03:00) Buenos Aires
 (GMT-03:00) Georgetown
 (GMT-03:00) Greenland
 (GMT-03:00) Montevideo
 (GMT-02:00) Mid-Atlantic
 (GMT-01:00) Azores
 (GMT-01:00) Cape Verde Is.
 (GMT) Casablanca
 (GMT) Greenwich Mean Time : Dublin, Edinburgh, Lisbon, London
 (GMT) Monrovia, Reykjavik
 (GMT+01:00) Amsterdam, Berlin, Bern, Rome, Stockholm, Vienna
 (GMT+01:00) Belgrade, Bratislava, Budapest, Ljubljana, Prague
 (GMT+01:00) Brussels, Copenhagen, Madrid, Paris
 (GMT+01:00) Sarajevo, Skopje, Warsaw, Zagreb
 (GMT+01:00) West Central Africa
 (partial listing)

Log (System Log) Configuration

Configure System Logging (Syslog) on this page from **Configuration > System > Log**. The Syslog data is stored in LIB-44xx RAM by default. Syslog data will be lost with an LIB-44xx reboot unless other provisions are made to save it.

For syslog [monitoring](#) details, see 'Monitor > System > Log' on page 314.



Server Mode

Sets / indicates the server mode operation. When Server mode is enabled, the syslog message will be sent 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 to the sender since UDP is a connectionless protocol and it does not provide acknowledgments. The syslog packet will always be sent out, even if the syslog server does not exist. Possible modes are:

Enabled: Enable server mode operation.

Disabled: Disable server mode operation.

Server Address

Sets / indicates the IPv4 or IPv6 host address of the syslog server. If the switch provides the DNS feature, it also can be a host name. If you enter an invalid address, or do not enter any address, the message “*The format of Server Address is invalid*” displays. Re-enter a valid IPv4 or IPv6 server address.

Syslog Level

Sets / indicates what kind of message will be sent to the syslog server. Possible modes are:

Info: Send information, warnings, and errors (all syslog information available).

Warning: Send just warnings and errors.

Error: Send just errors.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Ports Configuration

This page displays current port configurations and allows LIB-44xx port configuration from the **Configuration > Ports** menu path.

The **LIB-4400** is a 4-port device with one MGMT port, one CONSOLE port, and four SFP interfaces which can be operated in Auto, 1G, or 10G mode.

The **LIB-4424** is a 28-port device with one MGMT port, one CONSOLE PORT, twenty-four 100/1000 SFP interfaces which can be operated in Auto, 1G, or 10G mode.

The LIB-44xx Ethernet ports are equipped with LEDs for visual status of speed, duplex and activity. The Ethernet ports provide standard features such as configuring Auto negotiation, speed, duplex, flow control. These features are compliant with the IEEE 802.3-2008 Ethernet PHYs standards. Advertisement capabilities and autocross are always on. Auto negotiation is only supported in Auto mode. 1Gbps FDX is a forced mode.

The SFP port when in 100BaseFx mode is set at 100Mbps and the duplex mode can be user configured.

The SFP port when in 1000BaseX mode always has Auto-negotiation and Auto-negotiation Bypass modes enabled. The SFP port can operate in SGMII mode with Auto-negotiation always on. Note that different port types support different options; refer to each port type for capabilities.

Configuration > Ports > Configuration

The LIB-44xx ports can be configured here in terms of speed, flow control, max. frame size, excessive collision control, power control, and port description.

Port	Link	Speed		Adv Duplex		Adv speed		Flow Control			PFC		Maximum Frame Size	Excessive Collision Mode	Frame Length Check
		Current	Configured	Fdx	Hdx	10M	100M	1G	Enable	Curr Rx	Curr Tx	Enable			
*		Down	Auto	<input checked="" type="checkbox"/>	0-7	10240	<->								
1	Down	Auto	<input checked="" type="checkbox"/>	0-7	10240	Discard									
2	Down	Auto	<input checked="" type="checkbox"/>	0-7	10240	Discard									
3	Down	Auto	<input checked="" type="checkbox"/>	0-7	10240	Discard									
4	Down	Auto	<input checked="" type="checkbox"/>	0-7	10240	Discard									
5	Down	Auto	<input checked="" type="checkbox"/>	0-7	10240	Discard									
6	Down	Auto	<input checked="" type="checkbox"/>	0-7	10240	Discard									
7	Down	Auto	<input checked="" type="checkbox"/>	0-7	10240	Discard									
8	Down	Auto	<input checked="" type="checkbox"/>	0-7	10240	Discard									
9	Down	Auto	<input checked="" type="checkbox"/>	0-7	10240	Discard									
10	Down	Auto	<input checked="" type="checkbox"/>	0-7	10240	Discard									
11	Down	Auto	<input checked="" type="checkbox"/>	0-7	10240	Discard									

The Port Configuration parameters are explained below:

Port

This is the logical port number for this row (e.g., 1-5 for the LIB-4400). The * in the Port column acts as a 'wild card' character which causes the selections in this row to be applied to all other Ports (rows) in the table for which this selection is valid. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Link

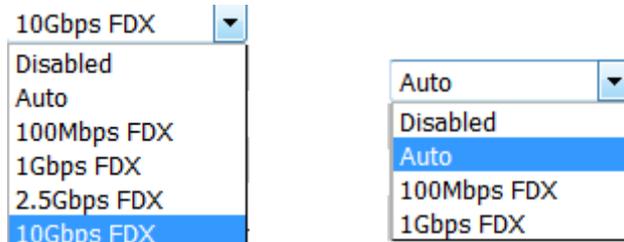
The current link state is displayed graphically. Green indicates the link is up and red that it is down.

Current Link Speed

Provides the current link speed of the port (e.g., **1Gfdx** for 1 Gbps/Full duplex, **100hdx** for 100 Mbps/Half duplex, or Down).

Configured Link Speed

Select any available link speed for the given LIB-44xx port. The available selections are **Disabled**, **Auto**, **100Mbps FDX**, **1Gbps FDX**, **2.5Gbps FDX** and **10Gbps FDX**.



Port 25 – 28 Config speed dropdown Port 1- 24 Config speed dropdown

<>: Displays in the 'wild card' row to indicate no parameter selection has been made yet.

Disabled: disables the LIB-44xx port operation (e.g., Ports 1-5 on the LIB-4400).

Auto: Port auto negotiates its speed with the link partner and selects the highest speed that is compatible with the link partner (e.g., Ports 1-5 on the LIB-4400). Both sides must be set the same. Depending on the type of port, (copper or SFP), this can be displayed as:

<>: Displays in the 'wild card' row to indicate no parameter selection has been made yet.

10Mbps HDX: Forces the port to 10 Mbps at Half duplex.

10Mbps FDX: Forces the port to 10 Mbps at Full duplex

100Mbps HDX: Forces the port to 100 Mbps at Half duplex.

100Mbps FDX: Forces the port to 100 Mbps at Full duplex.

1Gbps FDX: Forces the port to 1Gbps at Full duplex. Both sides must be set the same.

2.5Gbps FDX: Forces the port to 2.5Gbps at Full duplex. Both sides must be set the same.

10Gbps FDX: Forces the port to 10Gbps at Full duplex.

The 10/100/1000BaseT and 100BaseFx/1000BaseX/SGMII ports support auto-negotiation per the IEEE 802.3 standard.

The 10/100/1000BaseT also supports disabling Auto-negotiation and can be forced to 10 half, 10 full, 100 half, 100 full or 1G full-duplex modes. When auto negotiation parallel detects a forced mode remote, it defaults to the link speed and Half duplex. This can result in a forced Full duplex port talking to an Auto port operating in Half duplex mode, resulting in excessive collision issues. The Auto-negotiation signalling is compliant with IEEE802.3 2008 Clause 28.

In 1000BaseX Auto mode, auto negotiation can be enabled. In a case where the link partner doesn't auto negotiate, the bypass mode is activated automatically to link at 1000Mbps and full duplex. The Auto-negotiation signalling is compliant with IEEE802.3 2008 Clause 37.

In 100BaseFx mode, auto-negotiation is not supported. Duplex is configurable to half or full to support legacy equipment.

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 the speed that 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 checkbox to use flow control. This setting is related to the setting for Configured Link Speed. A green check mark (✓) indicates flow control is currently enabled for this port; a red x mark (✗) indicates flow control is currently disabled for this port.

Maximum Frame Size

Enter the maximum frame size to be allowed for the LIB-4400 port, including FCS. The valid range is 1522 - 10056 bytes. The default is 10056 bytes for 10GB SFP+ ports and 1522 for port 5.

Excessive Collision Mode

Configure port transmit collision behaviour:

<>: wild card character selects all.

Discard: Discard frame after 16 collisions (default).

Restart: Restart the backoff algorithm after 16 collisions.

The Excessive Collision Mode parameter applies to certain ports only (e.g., LIB-4400 port 5 only).

Power Control

The Usage column shows the current percentage of the power consumption per port. The 'Configured' column allows for changing the power savings mode parameters per port. The Power Control parameter applies only to certain ports

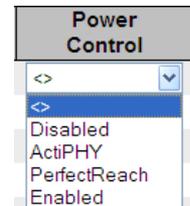
<>: wild card character selects all.

Disabled: All power savings mechanisms disabled.

ActiPHY: Link down power savings enabled. ActiPHY™ is an automatic power savings mode when a specific port is in link down or standby operation.

PerfectReach: Link up power savings enabled. PerfectReach™ is one of the LIB-4400 energy efficient modes where an intelligent algorithm actively determines the needed power level based on cable length.

Enabled: Both link up and link down power savings enabled.



Description

Lets you enter a definitive description for each port. The Port description string to uniquely identify the circuit. The description can be up to 31 alphanumeric characters.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Refresh: Click to refresh the page. Any changes made locally will be undone.

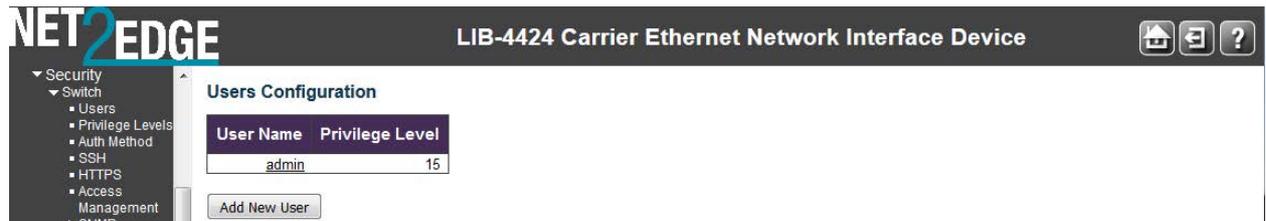
System Users Configuration

Configuration > Security > Switch > Users

The **Configuration > Security > Switch > Users** menu path lets you view and configure the system users that are allowed to access the web pages or log in from CLI.

You can also access the LIB-44xx System Password page from the **Configuration > Security > Switch > Users** menu path.

This page provides an overview of the currently defined system users. Currently the only way to login as another user on the web server is to close and reopen the browser.



User Name

Enter the new user's name to be added (the name identifying the user). This is also a link to Add, Edit or Delete an existing User.

Privilege Level

Enter the new user's level of access to be allowed. This is the privilege level of the user. The allowed range is **1** to **15**.

If the privilege level value is **15**, a user can access all groups (i.e., this user is granted the fully control of the device). But others value need 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.

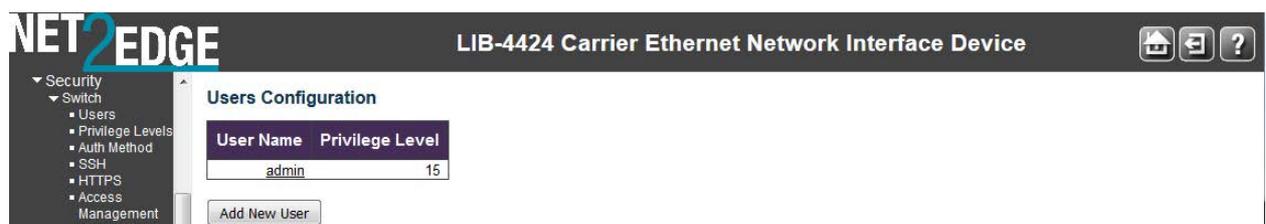
The system maintenance (software upload, factory defaults and etc.) requires user privilege level 15. By default, most groups' privilege level **5** has read-only access and privilege level **10** has read-write access. Generally, privilege level **15** can be used for an administrator account, privilege level **10** for a standard user account, and privilege level **5** for a guest account.

Buttons

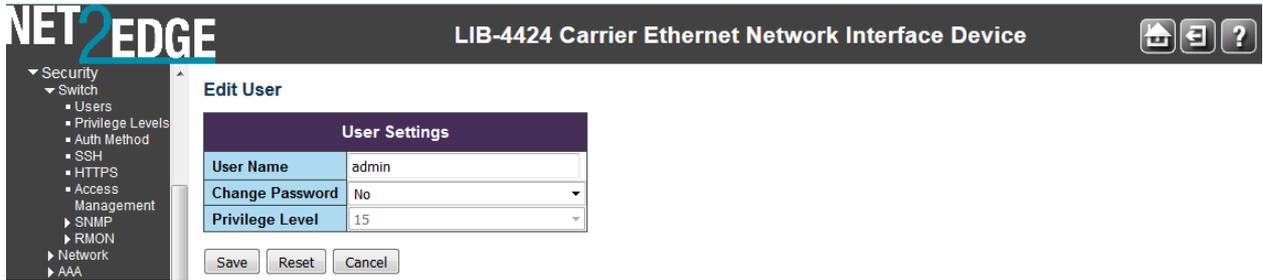
Add new user: Click to add a new user to the Users Configuration table.

Edit User (Edit the Default *Admin* User)

To edit the default admin user, click the [admin](#) link:



Edit the **Edit User** page.



Enter the **Password** and **Privilege Level** as described below:
Click the **Save** button when done.

User Name

A string identifying the user name that this entry should belong to. The allowed string length is **1** to **32**.

A valid user name can include a combination of letters, numbers and underscores.

Password

The password of the user. The allowed string length is **0** to **32** alpha, numeric, or special characters.

Password (again)

Enter the Password again to confirm. These entries must match exactly.

The new password must be entered twice to catch typing errors. The message “*Password Error - The old password is incorrect. New password is not set.*” displays if the new password entered is the same as the old password. If this occurs, click the browser Back button and enter a unique new password and confirm with an identical entry.

Privilege Level

The privilege level of the user. The allowed range is **1** 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 other values may be needed to refer to each group privilege level. A User's privilege level should be same or greater than the Group privilege level to have the access of that group.

In general, the user privilege levels are:

Privilege Level 15 can be used for an Administrator account. Privilege level 15 allows system Maintenance menu access (software upload, factory defaults, etc.).

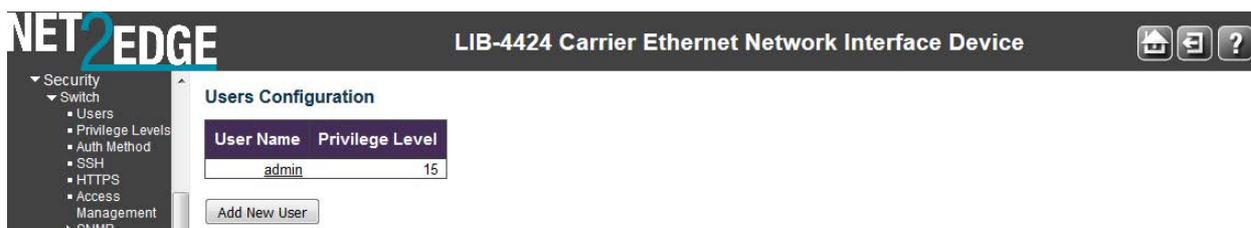
Privilege Level 10 for a Standard (basic) user account. Privilege level 10 allows read-write access.

Privilege Level 5 for a Guest account. By default, most groups privilege are assigned privilege level 5 (read-only access).

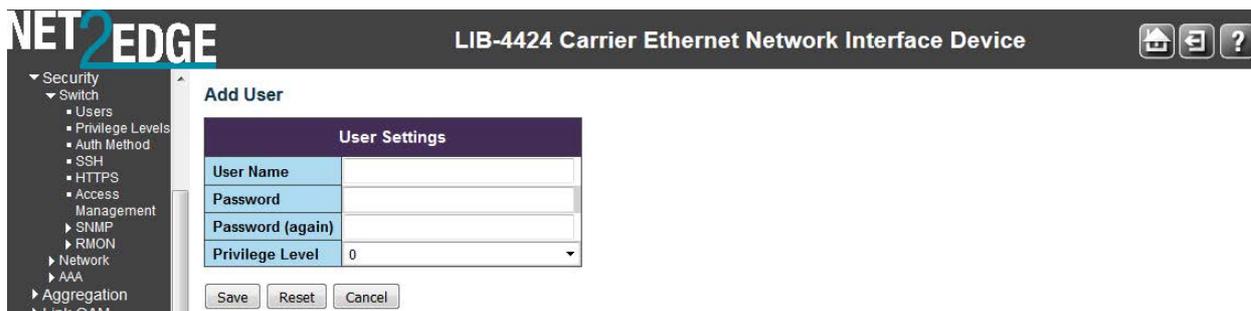
See the “Privilege Level Configuration” section below for more information.

Add a New User

To add a new user, click the **Add new user** button. The **Add User** table displays.

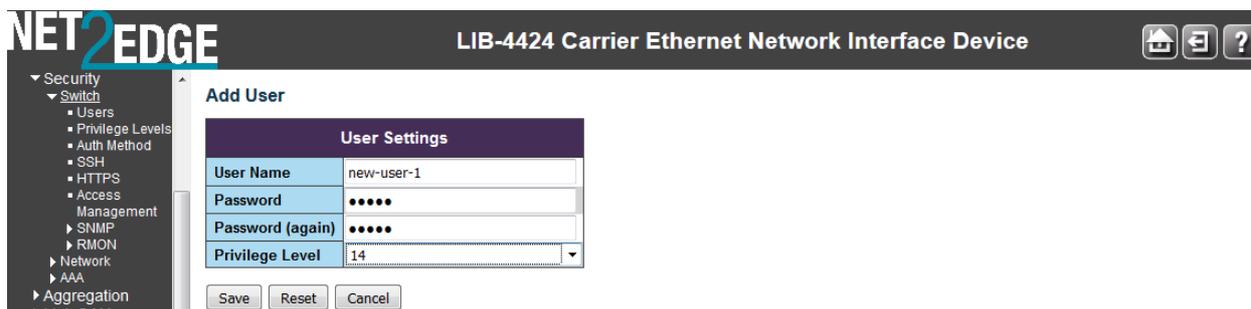


Enter the **User Name**, **Password**, and **Privilege Level** as described below:
Click the **Save** button when done.



Edit an Existing User's Parameters

To edit the default admin user, click the user's name (e.g., click **admin**). The **Edit User** table displays.



This page lets you configure the system password required to access the web pages or log in from the CLI. The parameters are explained below:

User Name

A string identifying the user name that this entry should belong to. The allowed string length is **1** to **32**.

A valid user name can include a combination of letters, numbers and underscores.

Password

The password of the user. The allowed string length is **0** to **32** alpha, numeric, or special characters.

Password (again)

Enter the Password again to confirm. These entries must match exactly.

The new password must be entered twice to catch typing errors. The message "*Password Error - The old password is incorrect. New password is not set.*" displays if the new password entered is the same as the old password. If this occurs, click the browser Back button and enter a unique new password and confirm with an identical entry.

Privilege Level

The privilege level of the user. The allowed range is **1** 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 other values may be needed to refer to each group privilege level. A User's privilege level should be same or greater than the Group privilege level to have the access of that group.

In general, the user privilege levels are:

Privilege Level 15 can be used for an Administrator account. Privilege level 15 allows system Maintenance menu access (software upload, factory defaults, etc.).

Privilege Level 10 for a Standard (basic) user account. Privilege level 10 allows read-write access.

Privilege Level 5 for a Guest account. By default, most groups privilege are assigned privilege level 5 (read-only access).

See the “Privilege Level Configuration” section below for more information.

Buttons

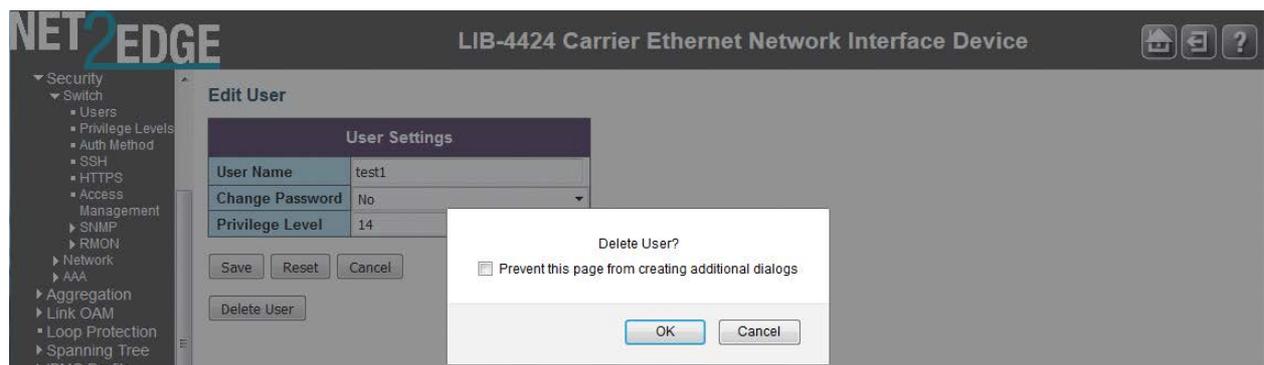
Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Cancel: Click to undo any changes made locally and return to the Users.

Delete an Existing User

To delete an existing user from the table, click the user's name to be deleted. The Edit User table displays.



Click the **Delete User** button. At the confirmation webpage message (*Delete User?*), click the **OK** button.

The Users Configuration table re-displays without the deleted user.

Privilege Level Configuration

Configuration ->Security > Switch - Privilege Levels

This page lets you view and edit users' privilege (access) levels from the **Configuration ->Security > Switch - Privilege Levels** menu path.

NET2EDGE LIB-4424 Carrier Ethernet Network Interface Device

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
Debug	15	15	15	15
DHCP	5	10	5	10
DHCPv6_Client	5	10	5	10
Diagnostics	5	10	5	10
EPS	5	10	5	10
ERPS	5	10	5	10
ETH_LINK_OAM	5	10	5	10
EVC	5	10	5	10
Firmware	5	10	5	10
Green_Ethernet	5	10	5	10
HQoS	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
MEP	5	10	5	10
Miscellaneous	15	15	15	15
MRP	5	10	5	10
MVR	5	10	5	10
MVRP	5	10	5	10
NTP	5	10	5	10
Performance_Monitor	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
RFC2544	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
TT_LOOP	5	10	5	10
UDLD	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
Y.1564(SAM)	5	10	5	10

The **Group Name** column lists the LIB-44xx main functions, including Aggregation, Diagnostics, EPS, ERPS, EtherSAT, ETH_LINK_OAM, EVC, IP, IPMC_LIB, IPMC Snooping, LACP, LLDP, Loop_Protect, MAC_Table, MEP, MVR, Maintenance, Mirroring, PHY, PTP, Port_Security, Ports, Private_VLANs, QoS, SNMP, Security, Spanning Tree, System, Timer, VCL, VLAN_Translation, and VLANs.

The Privilege Level parameters are explained below:

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 some groups contain more than one module.

Some of these privilege level groups are explained below:

System: e.g., Contact, Name, Location, Timezone, Log.

Security: Authentication, System Access Management, Port (contains Dot1x port, MAC based and the MAC Address Limit), ACL, HTTPS, SSH, ARP Inspection and IP source guard.

IP: Everything except 'ping'.

Port: Everything except 'VeriPHY'.

Diagnostics: 'ping' and 'VeriPHY'.

Maintenance: CLI- System Reboot, System Restore Default, System Password, Configuration Save, Configuration Load and Firmware Load. Web- Users, Privilege Levels and everything in Maintenance.

Privilege Levels

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 statistics). User Privilege should be same or greater than the authorization Privilege level to have the access to that group.

Configuration read-only: these users are only allowed to monitor status / configuration settings.

Configuration/execute read-write: these users are only allowed to monitor status and make changes to configuration settings.

Status/statistics read-only: these users are only allowed to monitor status / statistics settings.

Status/statistics read-write: (e.g., for clearing statistics).

User Privilege Levels (1-15)

The privilege level of the user. The allowed range is **1** to **15**.

If the privilege level value is 15, it can access all groups (i.e. it is granted full control of the device). But other values need to refer to each group privilege level. User's privilege should be same or greater than the group privilege level in order to have the access of that group. By default, most groups have privilege level 5 with read-only access; privilege level 10 has the read-write access. The system maintenance functions (software upload, factory defaults, etc.) require user privilege level 15. Generally, the user privilege levels are:

Privilege Level 15 can be used for an Administrator account,

Privilege Level 10 is for a Standard (basic) user account, and

Privilege Level 5 is for a Guest account.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Message: *Must explicitly add case for module ID: %d*

Meaning: A user privilege level was not defined for a module.

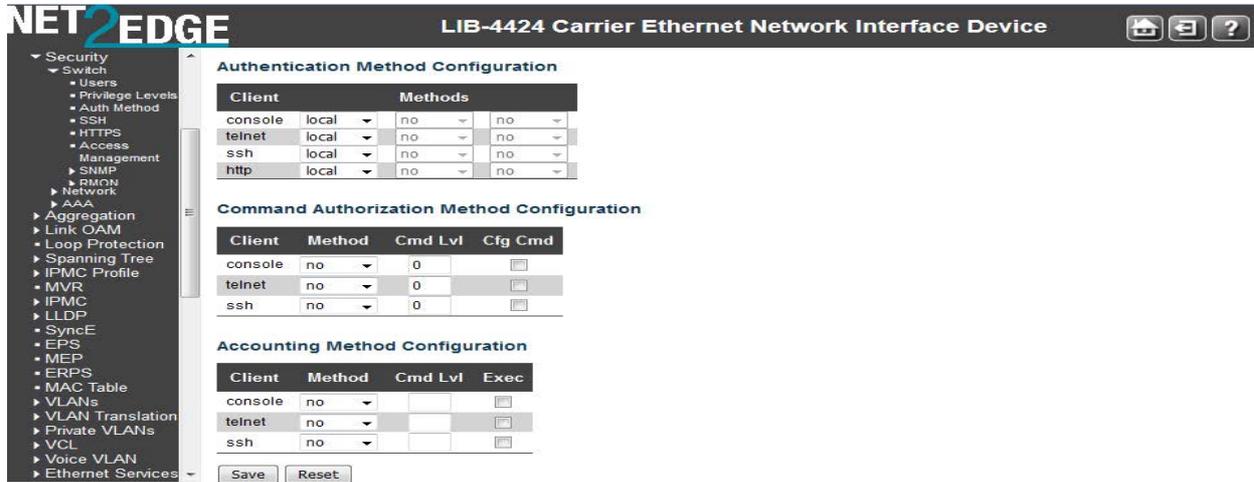
Recovery:

1. Define a privilege level for the module (e.g., crw = 10).

Authentication Method Configuration

Configuration > Security > Switch > Auth Method

This page lets you configure how a user is authenticated when they log into the LIB-44xx via one of the management client interfaces. Access this page from the **Configuration > Security > Switch > Auth Method** menu path.



The table has one row for each client type and several columns, as explained below:

Client

The management client for which the configuration below applies (console, telnet, ssh, web).

Authentication Method

Authentication Method can be set to one of the following values:

none: authentication is disabled and login is not possible via this method (see **Note** below).

local: use the local user database on the LIB-44xx for authentication.

RADIUS: use a remote RADIUS server for authentication.

TACACS+: use a remote TACACS+ server for authentication.

Note: setting the Authentication Method to **none** for all four methods disallows subsequent login. Do **not** set all of the management client authentication methods to '**none**' as this will disable login.

Fallback

Enable fallback to local authentication by checking this checkbox.

If none of the configured authentication servers are alive, the local user database is used for authentication. This is only possible if the 'Authentication Method' is set to a value other than 'none' or 'local'.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

SSH Configuration

Configure SSH on this page from the **Configuration - Security - Switch - SSH** menu path. The Secure Shell (SSH) network protocol allows data to be exchanged using a secure channel between

two networked devices. The encryption used by SSH provides confidentiality and integrity of data over an unsecure network.

The LIB-44xx CLI can be accessed over the SSH interface. This is provided for customer who need a more secure interface. SSH uses public-key cryptography for authentication. When the SSH server is enabled, normal telnet access can be enabled or disabled to avoid any security holes.

All SSH parameters and certificates uploaded to the LIB-44xx are available on power cycle. The LIB-44xx ships with a default certificate; it is the end user's responsibility to get a valid certificate from a certificate authority (e.g., from Verisign, DigiCert, Thawte, etc.).

See “[HTTPS Configuration](#)” on page 37 for certificate upload information.



Mode

Indicates the SSH mode operation. Possible modes are:

Enabled: Enables SSH mode operation.

Disabled: Disables SSH mode operation.

Buttons

Save: Click to save changes.

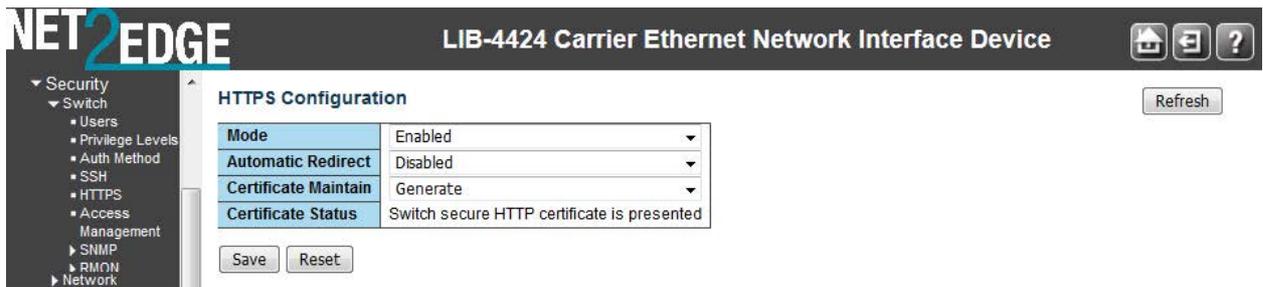
Reset: Click to undo any changes made locally and revert to previously saved values.

HTTPS Configuration

Configure HTTPS on this page from the **Configuration - Security - Switch - HTTPS** menu path. Hypertext Transfer Protocol over Secure Socket Layer (HTTPS) is used to indicate a secure HTTP connection. HTTPS provides authentication and encrypted communication.

The LIB-44xx has an embedded web server for managing the device without any additional software. The web server also provides a secure interface using HTTPS. The validity period will be based on the validity period of the uploaded cert. If using a generated cert, then the HTTPS Certificate's validity period is from Jan. 1, 2010 to Dec. 31, 2029.

By default the servers listen on standard port 80 for HTTP and on standard port 443 for HTTPS. The HTTPS runs over SSL and the certificate can be uploaded by the user using standard TFTP protocol. You can reconfigure the HTTPS server port for security purposes. No password is used for the SSH certificate. Open SSL commands are available for self-signing a certificate.



HTTPS Configuration

Mode

Indicates / sets the HTTPS mode operation. The possible modes are:

Enabled: Enable HTTPS mode operation. After a change from 'Disabled' to 'Enabled' and a 'Save', you must login in secure mode (i.e., from <https://192.168.1.110>).

Disabled: Disable HTTPS mode operation. After a change from 'Enabled' to 'Disabled' and a 'Save', you must login in non-secure mode (i.e., from <http://192.168.1.110>).

Automatic Redirect

Indicates the HTTPS redirect mode operation. Automatically redirect web browser to HTTPS when HTTPS mode is enabled. Possible modes are:

Enabled: Enable HTTPS redirect mode operation.

Disabled: Disable HTTPS redirect mode operation. Note: You cannot enable the HTTPS redirect function when the HTTPS operation mode is disabled.

The valid HTTPS configurations are:

Mode = Disabled and **Automatic Redirect** = Disabled, or

Mode = Enabled and **Automatic Redirect** = Disabled, or

Mode = Enabled and **Automatic Redirect** = Enabled.

If you enable both "Mode" and "Automatic Redirect", these messages display: "Content was blocked because it was not signed by a valid security certificate." and "There is a problem with this website's security certificate.". Select "Continue to this website (not recommended)". At the login dialog box, enter the login information. The startup screen ("Port State Overview") displays from the new (secure) login IP address (e.g., <https://192.168.1.110/>).

HTTPS Certificate

View

Displays the current HTTPS certificate (see above). Each certificate contains Data and a Signature Algorithm. A sample HTTPS Certificate View display is shown below:

Certificate:

Data:

Version: 3 (0x2)

Serial Number: 9 (0x9)

Signature Algorithm: sha1WithRSAEncryption

Issuer: CN=LIB-4400

Validity

Not Before: Jan 1 00:00:00 2010 GMT

Not After : Dec 31 00:00:00 2029 GMT

Subject: CN=LIB-4400

Subject Public Key Info:

Public Key Algorithm: rsaEncryption

RSA Public Key: (512 bit)

Modulus (512 bit):

```
00:d6:78:4a:22:4f:9c:8c:83:1f:8e:c0:65:d6:db:
6f:ca:40:91:c3:67:39:08:d4:e0:c6:06:81:a4:81:
a4:da:ca:52:1e:95:dc:18:1f:d3:5e:47:2f:d9:4d:
1e:76:37:ac:b4:cb:b9:36:d1:32:19:69:3a:19:0b:
a9:33:5e:e1:2b
```

Exponent: 65537 (0x10001)

Signature Algorithm: sha1WithRSAEncryption

```
14:2f:40:fd:eb:9f:0e:ac:e6:9a:2e:8c:ff:70:65:a6:19:7b:
1c:50:15:7e:01:6c:46:5e:ba:8c:77:ae:b6:20:40:a8:96:1c:
41:a1:fe:d2:0a:16:57:59:c2:5f:74:4b:82:6c:f5:4f:57:5d:
a3:e5:94:1b:da:e0:43:c1:49:1b
```

Depending on the encryption method selected (RSA or DSA) below, the **HTTPS Certificate > View** section displays one of two signature algorithms in effect, either:

Signature Algorithm: dsaWithSHA1 or

Signature Algorithm: sha1WithRSAEncryption

Generate

The Generate dropdown lets you select **RSA** or **DSA** as the public key algorithm.

RSA: generate an RSA key. Uses the RSA internet encryption and authentication system via an algorithm developed in 1977 by Ron Rivest, Adi Shamir, and Leonard Adleman.

DSA: generate a DSA key. Uses the digital signature algorithm (DSA) developed by the U.S. National Security Agency (NSA) to generate a digital signature for the authentication of electronic documents. Digital signatures are generated and verified via DSA. Signatures are generated in conjunction with the use of a private key; verification takes place in reference to a corresponding public key.

When done, click the **Generate** button to generate the certificate as defined at the Generate dropdown.

Depending on the encryption method selected (RSA or DSA) here, the **HTTPS Certificate > View** section displays one of two signature algorithms in effect, either:

Signature Algorithm: dsaWithSHA1 or

Signature Algorithm: sha1WithRSAEncryption

Load

Lets you browse to and select a new HTTPS certificate and load it.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Browse...: Click to display the **Choose a File to Upload** dialog box. Select a file and click **Open**.

Load: Click to load the selected certificate file.

Generate a Self-signed Certificate

A self signed certificate can be generated using openssl.

1. Type **openssl genpkey -algorithm RSA -pkeyopt rsa_keygen_bits:512 -out LIB-4400-key.pem**.

2. Type **openssl req -key LIB-4400-key.pem -x509 -new -set_serial 2 -days 3650 -out LIB-4400-cacert.pem.**
3. Type **cat LIB-4400-cacert.pem LIB-4400-key.pem > LIB-4400-cacert-key.pem.**

The extension is .pem by convention. The Serial Number is hard coded to 2 (0x2).

For more OpenSSL information, see the “openssl command line tool” web page at <http://www.openssl.org/docs/apps/openssl.html>.

Continue with the “Load a Certificate from CLI” procedure below:

Load a Certificate from the CLI

1. Start and configure the TFTP server.
2. Place the certificate to load in your tftpboot root (e.g., *C:\TFTP-Root*) and run:

Security Switch HTTPS Certificate Load <ip-of-tftp-server> <cacert-key.pem>

An example of a valid certificate is shown below:

```
-----BEGIN CERTIFICATE-----
MIICKTCCAjugAwIBAgIBHjANBgkqhkiG9w0BAQUFADCBpzELMAkGA1UEBhMCVVMx
EjAQBgNVBAgMCU1pbm5lc290YTETMBEGA1UEBwwKTWlubmV0b25rYTEhMB8GA1UE
CgwYVHJhbnNpdGlvbiBOZXR3b3JrcywgSW5jMRQwEgYDVQQLDAtFbmdpbmVlcmlu
ZzEVMBMGA1UEAwwMMMTkyLjE2OC4xLjIwMR8wHQYJKoZIhvcNAQkBFhB1c2VyQGV4
YW1wbGUuY29tMB4XDTEyMTAzMTEzNi0xMTAyODE3MTIzNlowgacxZzAJ
BgNVBAYTAiVMTMlEAYDVQQIDAINaW5uZXRvdGEExEzARBgNVBAcMCK1pbm5ldG9u
a2ExITAfBgNVBAoMGFRyYW5zaXRpb24gTmV0d29ya3MslEluYzEUMBIGA1UECwwL
RW5naW5lZXJpbmxcFTATBgNVBAMMDDE5Mi4xNjguMS4yMDEfMBOGCSqGSIb3DQEJ
ARYQdXNlckBleGFtcGxILmNvbTBcMA0GCSqGSIb3DQEBAQUAA0sAMEgCQQC9HslP
pozHOALDNNiZlgSdOPpzD9SB+ftfoxGfOao4abrHpp4HAUesB7vkvVy0mWztngn
pOHanBmYmsP1L/YBAgMBAAGjUDBOMB0GA1UdDgQWBHTHTvxNJ7ycaKsB6NyBM62s
O8svIDAfBgNVHSMEGDAWgBTHTvxNJ7ycaKsB6NyBM62sO8svIDAMBgNVHRMEBTAD
AQH/MA0GCSqGSIb3DQEBBQUAA0EAmiLPKdgZ5sWaVKwznqZOZWWZsqF6FspCb9Sn4
RHR8vltDwnE//2sWt73XZjhK90w5oe+3CG9j0I0H6IU12MHK3Q==
-----END CERTIFICATE-----
-----BEGIN PRIVATE KEY-----
MIIBUwIBADANBgkqhkiG9w0BAQEFAASCAT0wggE5AgEAAkEAvR7JT6aMxzgCwzTY
mSIEnTj6cw/Ugfn7X6MRnzmqOGm64R6aeBwFHRaE75L1ctJls7Z4J6Th2jWzMprD
9S/2AQIDAQABAKB3CzsC9T6mZF5wzAIPmnc3FFeqN+u3iUivMbf0/dJIQSSbkuYA
73BJxKXczfcsA74j78wxWdMMZUt6Zy3y2I0AiEA6MbzaQ6SZhlyEFdAiaA8P9mY
4Yf6G3lyHro2VdRw2qsCIQDP/Np3hJ9QiW9IPggKeBjLKLXC8w7WhvDL3qApX3gy
AwlgY/PR30GXkWnJtcSRfMcmN9CKugT2W0Ssys7wt+gx0WMCIG3amK1d8CPCo3Nd
yDN3YWkUtlQmiuTGWzLjAGdwwsxvAiBq/xNcMvLaQgJSeQmTQmSZvSRKYsl/TjRr
FpU5omhpSQ==
-----END PRIVATE KEY-----
```

Note that a valid certificate consists of both a Certificate and a Private Key.

The SSH / RSA parameters are:

SSH Certificate Thread Stack Size:	65536 bytes
SSH Management Max. Host Key Length:	4096 bytes

Host Key Length: 4096 bytes
 Max. RSA Size: 1024 bit

Access Management Configuration

Configure access management table on this page from **Configuration - Security - Switch - Access Management**. You can add up to 16 entries. If the application's type matches any one of the access management entries, it will allow access to the LIB-44xx.

Click the **Add New Entry** button to start configuring a new Access Management entry.

NET2EDGE LIB-4424 Carrier Ethernet Network Interface Device

Access Management Configuration

Mode: Disabled

Delete	VLAN ID	Start IP Address	End IP Address	HTTP/HTTPS	SNMP	TELNET/SSH
Delete	1	0.0.0.0	0.0.0.0	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Add New Entry

Save Reset

Mode

Indicates / sets the access management mode operation. The possible modes are:

Enabled: Enable access management mode operation.

Disabled: Disable access management mode operation.

Delete

Click to delete the entry. It will be deleted during the next save.

Start IP Address

Indicates / sets the beginning IP address for the access management entry. The value of Start IP Address must be a valid IP address in dotted decimal notation ('x.y.z.w'). The following restrictions apply:

- 1) x, y, z, and w must be decimal numbers from **0** - **255**,
- 2) x must not be **0**,
- 3) x must not be **127**, and
- 4) x must not be greater than **223**.

End IP Address

Indicates / sets the ending IP address for the access management entry. The value of End IP Address must be a valid IP address in dotted decimal notation ('x.y.z.w'). The following restrictions apply:

- 1) x, y, z, and w must be decimal numbers from **0** - **255**,
- 2) x must not be **0**,
- 3) x must not be **127**, and
- 4) x must not be greater than **223**.

HTTP/HTTPS

Check to allow the host access to the LIB-44xx from the HTTP/HTTPS interface if the host IP address matches the IP address range provided in the entry.

SNMP

Check to allow the host access to the LIB-44xx from the SNMP interface if the host IP address matches the IP address range provided in the entry.

TELNET/SSH

Check to allow the host access to the LIB-44xx from the Telnet/SSH interface if the host IP address matches the IP address range provided in the Start IP Address / End IP Address entry (above).

Buttons

Add New Entry: Click to add a new access management entry.

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

SNMP Configuration

Configure SNMP on this page from the **Configuration > Security > Switch > SNMP > System** menu path.

Here you can configure LIB-44xx SNMP System, Communities, Users, Groups, Views, Access and trap parameters. Simple Network Management Protocol (SNMP) is part of the TCP/IP protocol for network management. SNMP allows diverse network objects to participate in a network management architecture. It enables network management systems to learn network problems by receiving traps or change notices from network devices running SNMP.

The SNMP agent embedded in the LIB-44xx is capable of version 1, 2c, or v3 support to access all management information from the device. The community strings for v1 and v2c and the USM/VACM for SNMPv3 are supported. The SNMP agent can support IPv4 and IPv6 trap destinations. It also supports the INFORM PDU for notification along with traps.

Traps are generated when a condition has been met on the SNMP agent. These conditions are defined in the Management Information Base (MIB). The administrator then defines thresholds, or limits to the conditions, that are to generate a trap. Conditions range from preset thresholds to a restart.

All of the values that SNMP reports are dynamic. The information needed to get the specified values that SNMP reports is stored in the MIB. This information includes Object IDs (OIDs), Protocol Data Units (PDUs), etc. The MIBs must be located at both the agent and the manager to work effectively.

SNMP v1, v2c, v3 Descriptions

Each SNMP version is described below:

SNMPv1

SNMP version 1 (SNMPv1) is the initial implementation of the SNMP protocol. SNMPv1 operates over protocols such as User Datagram Protocol (UDP), Internet Protocol (IP), OSI Connectionless Network Service (CLNS), AppleTalk Datagram-Delivery Protocol (DDP), and Novell Internet Packet Exchange (IPX). SNMPv1 is widely used and is the de facto network-management protocol in the Internet community. The first RFCs for SNMP, now known as SNMPv1, appeared in 1988: RFC 1065, RFC 1066, and RFC 1067. These protocols were obsoleted by SNMPv1: RFC 1155, RFC 1156 and RFC 1157. After a short time, RFC 1156 (MIB-1) was replaced by the more often used *RFC 1213 - Version 2 of management information base (MIB-2) for network management of TCP/IP-based*

internets. SNMPv1 was criticized for its poor security. Authentication of clients is performed only by a "community string", in effect a type of password, which is transmitted in cleartext.

SNMPv2 and v2c

SNMPv2 (RFC 1441–RFC 1452) revises SNMPv1 and includes improvements in the areas of performance, security, confidentiality, and manager-to-manager communications. It introduced GetBulkRequest, an alternative to iterative GetNextRequests for retrieving large amounts of management data in a single request. However, the new party-based security system in SNMPv2, viewed by many as overly complex, was not widely accepted.

Community-Based Simple Network Management Protocol version 2, or SNMPv2c, is defined in RFC 1901–RFC 1908. In its initial stages, this was also informally known as SNMPv1.5. SNMPv2c comprises SNMPv2 without the controversial new SNMP v2 security model, using instead the simple community-based security scheme of SNMPv1. While officially only a "Draft Standard", this is widely considered the de facto SNMPv2 standard.

User-Based Simple Network Management Protocol version 2, or SNMPv2u, is defined in RFC 1909–RFC 1910. This is a compromise that attempts to offer greater security than SNMPv1, but without incurring the high complexity of SNMPv2. A variant of this was commercialized as SNMP v2*, and the mechanism was eventually adopted as one of two security frameworks in SNMP v3.

SNMPv3

Simple Network Management Protocol Version 3 (SNMPv3) is an interoperable standards-based protocol for network management. SNMPv3 provides secure access to devices by a combination of authenticating and encrypting packets over the network. Although SNMPv3 makes no changes to the protocol aside from the addition of cryptographic security, its developers have managed to make things look much different by introducing new textual conventions, concepts, and terminology.

SNMPv3 primarily added security and remote configuration enhancements to SNMP. Security has been the biggest weakness of SNMP since the beginning. Authentication in SNMP Versions 1 and 2 amounts to nothing more than a password (community string) sent in clear text between a manager and agent. Each SNMPv3 message contains security parameters which are encoded as an octet string. The meaning of these security parameters depends on the security model being used.

SNMPv3 provides important security features:

Confidentiality - Encryption of packets to prevent snooping by an unauthorized source.

Integrity - Message integrity to ensure that a packet has not been tampered with in transit including an optional packet replay protection mechanism.

Authentication - used to verify that the message is from a valid source.

SNMPv3 provides for both security models and security levels. A security model is an authentication strategy that is set up for a user and the group in which the user resides. A security level is the permitted level of security within a security model. A combined security model / security level determines which security mechanism is used when handling an SNMP packet. Three security models are available: SNMPv1, v2c, and v3.

SNMPv3 introduces the following key features:

SNMPv3 EngineID

SNMPv3 USM (User-Based Security Model)

SNMP VACM (View-based Access Control Model)

SNMP Trap/Inform (v1/v2c/v3 trap, v2c/v3 inform)

The SNMPv3 function supports these services:

SNMP v3 user management, authentication and encryption
 SNMP VACM management
 SNMP v1/v2c/v3 selection
 SNMP notification (v1/v2c/v3 trap, v2c/v3 inform) functionality

SNMP v3 EngineID concept: an SNMP engine provides services for sending and receiving messages, authenticating and encrypting messages, and controlling access to managed objects. There is a one-to-one association between an SNMP engine and the SNMP entity which contains it. The SNMP v3 engine contains a Dispatcher, a Message Processing Subsystem, a Security Subsystem, and an Access Control Subsystem.

Within an administrative domain, an `snmpEngineID` is the unique and unambiguous identifier of an SNMP engine. Since there is a one-to-one association between SNMP engines and SNMP entities, it also uniquely and unambiguously identifies the SNMP entity within that administrative domain. Note that it is possible for SNMP entities in different administrative domains to have the same value for `snmpEngineID`. Federation of administrative domains may necessitate assignment of new values.

SNMPv3 USM: the SNMP v3 implementation uses the traditional concept of a user (identified by a `userName`) with associated security information. This is a key SNMPv3 security feature implemented per RFC 3414.

SNMP v3USM User: Management operations using this Security Model make use of a defined set of user identities. For any user on whose behalf management operations are authorized at a particular SNMP engine, that SNMP engine must have knowledge of that user. An SNMP engine that wishes to communicate with another SNMP engine must also have knowledge of a user known to that engine, including knowledge of the applicable attributes of that user.

SNMPv3 VACM: The View-based Access Control Model defines a set of services that an application (such as a Command Responder or a Notification Originator application) can use for checking access rights. Access Control occurs (either implicitly or explicitly) in an SNMP entity when processing SNMP retrieval or modification request messages from an SNMP entity. Access Control also occurs in an SNMP entity when an SNMP notification message is generated (by a Notification Originator application).

VACM includes these elements: Groups, Security Levels, Contexts, MIB Views and View Families, and Access Policy.

SNMPv3 VACM – Groups: a Group is a set of zero or more `<securityModel, securityName>` tuples on whose behalf SNMP management objects can be accessed. A Group defines the access rights afforded to all `securityNames` which belong to that group. The combination of a `securityModel` and a `securityName` maps to at most one Group. A Group is identified by a `groupName`.

The Access Control module assumes that the `securityName` has already been authenticated as needed and provides no further authentication of its own. The View-based Access Control Model uses the `securityModel` and the `securityName` as inputs to the Access Control module when called to check for access rights. It determines the `groupName` as a function of `securityModel` and `securityName`.

Note that when the security model is v1 or v2c, the groups "public" and "private" cannot be removed, but when the security model is v3 the groups "public" and "private" can be removed.

SNMPv3 VACM – Views: Views are used to restrict the access rights of some groups to only a subset of the management information in the management domain.

A view subtree is the set of all MIB object instances which have a common ASN.1 OBJECT IDENTIFIER prefix to their names.

A family of view subtrees is a pairing of an OBJECT IDENTIFIER value (called the family name) with a bit string value (called the family mask). The family mask indicates which sub-identifiers of the associated family name are significant to the family's definition.

SNMPv3 Traps and Informs: A Trap is an SNMP message sent from one application to another (which is typically on a remote host). Their purpose is merely to notify the other application that something has happened, has been noticed, etc. The big problem with Traps is that they're unacknowledged, so you don't actually know if the remote application received your -important message. The trap is available for SNMP v1, v2c and v3.

An Inform is an acknowledged Trap. When the remote application receives an inform it sends back an acknowledgement message. Inform is available for SNMP v2c and v3. For SNMP v3, an inform must be sent to a specific remote USM user resided in the inform receiver.

SNMP v3 Users, Groups, and Views Configuration

SNMP v3 configuration involves setting up SNMP v3 Users, Groups, and Views, with the caveats discussed below: With SNMPv3, you can define SNMP users, groups, and views to provide access control to SNMP devices, and restrict certain users so they can only access the parts of the MIB that they have been given access rights to. The SNMP views, groups and users are described below:

Note: the concept of SNMP communities that was introduced in SNMPv2 is not relevant to SNMPv3, and has been replaced by SNMP groups/users. However, you can configure an LIB-44xx NID to respond to both SNMPv2 and SNMPv3 commands. If both SNMPv2 and SNMPv3 are to be used, you must configure SNMP communities and SNMP groups and users. To use SNMPv1/v2c, all you need to do is add a Community; nothing else needs to be configured. You can add or delete any Communities including the default communities.

Summary: You can create multiple Views. You can then create multiple Groups, and associate them with a View. You can configure multiple Groups (each with a different Group name and security level) and associate them with a particular View. You can also configure more than one View associated with a Group (e.g., a Group with read access to the entire MIB tree, but with only write access to certain objects). You could then create multiple Users, and associate them with a Group (you can associate multiple Users with a particular Group).

With SNMPv3 you can define SNMP views, groups and users to provide access control to SNMP devices, as shown below:

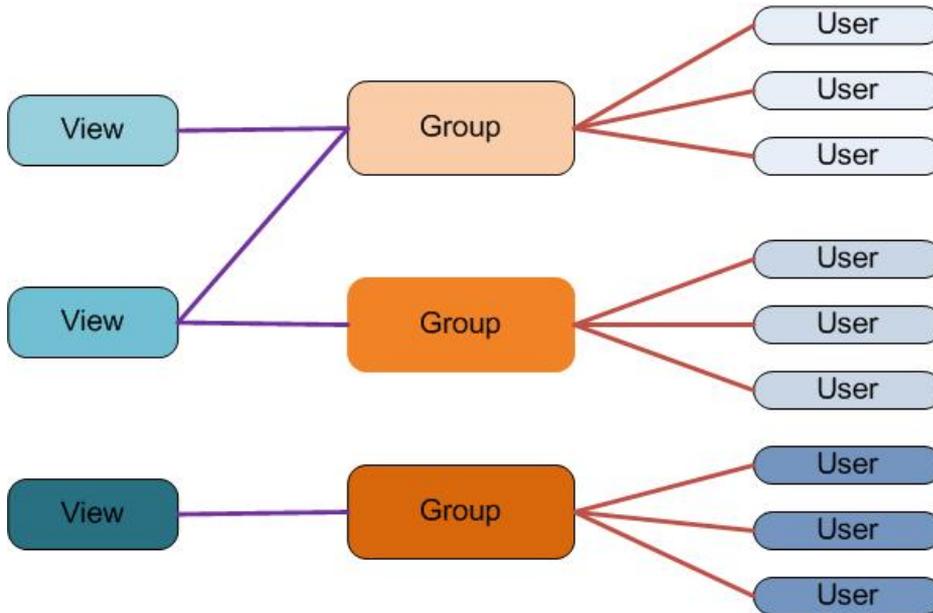


Figure 1. SNMP v3 Users, Groups, and Views

You can create multiple Views. You can then create multiple Groups, and associate them with a View. You can configure multiple Groups (each with a different Group name and security level) and associate them with a particular View. You can also configure more than one View associated with a Group (e.g., a Group with read access to the entire MIB tree, but with only write access to certain objects). You could then create multiple Users, and associate them with a Group (you can associate multiple Users with a particular Group).

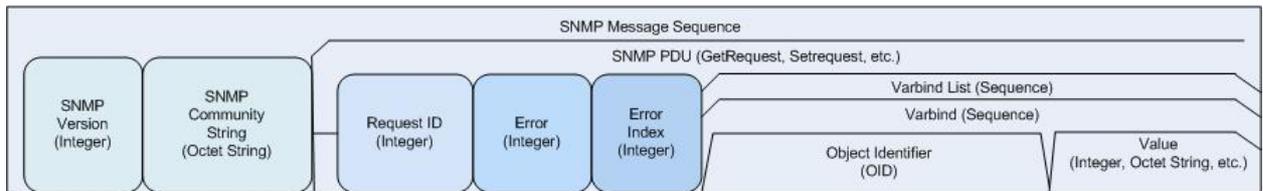
SNMP v1 Traps

Most SNMPv1 messages follow a model where a client (Network Management System) makes a request and a server (agent) responds to that request. Traps are the exception. An SNMP agent will transmit a trap to the NMS when it has a condition to report that is deemed too important to wait until asked. A common example of this is the failure of a communications link.

With most traps, the agent will include something called "'Interesting' variable bindings," which are the OID(s) and value(s) of MIB variable(s) that provide more information about the condition. So, for example, when a communications channel fails, the agent will send a pSError trap, which will have the OID of the "Link" or "Signal Detect" variable for that channel, and the value "down." Newer versions of Management Module firmware will also include bindings for the BIA (Cabinet serial number), slot, and (if applicable) subdevice of the entity in question. This information is always embedded in the first binding (as above), but are repeated separately for more convenient viewing under certain NMS packages.

SNMP v2 Traps

All LIB-44xx SNMP Trap messages conform to SNMPv2 MIB RFC-2573. See the "Supported MIBs" section below for information on support for public (standard) and private MIBs. A sample SNMP Message sequence is shown below:



SNMP v3 Traps

The SNMP v3 traps are mainly SNMPv2 traps with added authentication and privacy capabilities. SNMPv3 Traps use the engineID of the local application sending the trap rather than the engineID of the remote application. This means that you must create users in your remote user database and create one for each engineID you wish to send traps from.

See “[Appendix E: SNMP MIBs and Traps](#)” on page 518 for the Trap MIB variables.

SNMP Trap configuration is done at **Configuration > Security > Switch > SNMP > System**. See [SNMP System Configuration](#) on page 46.

SNMP v3 Configuration Process

Perform these procedures to configure the LIB-44xx for SNMP v3.

System - see [SNMP System Configuration](#) on page 46.

Communities - see [SNMPv3 Community Configuration](#) on page 49.

Users - see [SNMPv3 User Configuration](#) on page 50.

Groups - see [SNMPv3 Group Configuration](#) on page 52.

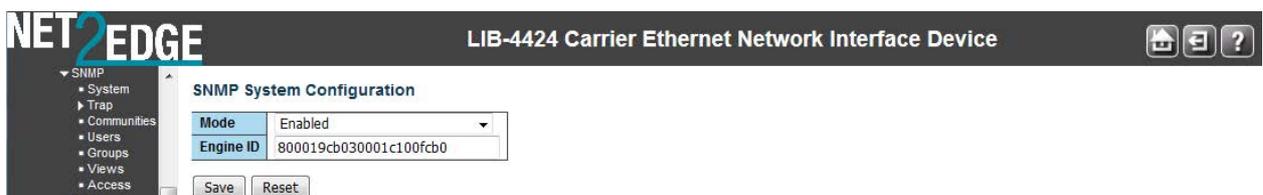
Views - see [SNMPv3 View Configuration](#) on page 53.

Access - see [SNMPv3 Access Configuration](#) on page 54.

The procedures in this process are defined in the following sections.

SNMP System Configuration

The default **Configuration > Security > Switch > SNMP > System** page is shown below: Note that the SNMP parameters displayed will vary depending on the SNMP Trap version selected.



SNMP Trap Configuration	
Trap Config Name	
Trap Mode	Disabled
Trap Version	SNMP v2c
Trap Community	public
Trap Destination Address	
Trap Destination Port	162
Trap Inform Mode	Disabled
Trap Inform Timeout (seconds)	3
Trap Inform Retry Times	5
Trap Security Engine ID	800019cb030001c100fcb0
Trap Security Name	None

Save Reset

All of the **Configuration > Security > Switch > SNMP > System** parameters are explained below:

Mode

Sets the SNMP mode operation. Possible modes are:

Enabled: Enable SNMP mode operation (default).

Disabled: Disable SNMP mode operation.

Version

Sets the SNMP supported version. Possible versions are:

SNMP v1: Set SNMP support to version 1.

SNMP v2c: Set SNMP support to version 2c (default).

SNMP v3: Set SNMP support to version 3.

Read Community

Indicates the community read access string to permit access to SNMP agent. The allowed string length is 0 to 255, and the allowed content is the ASCII characters from 33 to 126. The default is 'public'.

The field is applicable only when SNMP version is SNMPv1 or SNMPv2c. If SNMP version is SNMPv3, the community string will be associated with SNMPv3 communities table. It provides more flexibility to configure security name than a SNMPv1 or SNMPv2c community string. In addition to community string, a particular range of source addresses can be used to restrict source subnet. This field is greyed out if 'SNMP v3' is selected.

Write Community

Indicates the community write access string to permit access to SNMP agent. The allowed string length is 0 to 255, and the allowed content is the ASCII characters from 33 to 126. The default is 'private'.

The field is applicable only when SNMP version is SNMPv1 or SNMPv2c. If SNMP version is SNMPv3, the community string will be associated with SNMPv3 communities table. It provides more flexibility to configure security name than a SNMPv1 or SNMPv2c community string. In addition to community string, a particular range of source addresses can be used to restrict source subnet. This field is greyed out if 'SNMP v3' is selected.

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 (e.g.,

800007e5017f000001). Changing the Engine ID clears (deletes) all original local users. This field is greyed out unless 'SNMP V3' is selected.

SNMP Trap Configuration

Configure SNMP traps in this section from the **Configuration > Security > Switch > SNMP > System** menu path.

Trap Mode

Sets the SNMP trap mode operation. The valid selections are:

Enabled: Enable SNMP trap mode operation.

Disabled: Disable SNMP trap mode operation (default).

Trap Version

Sets the SNMP trap supported version. The valid selections are:

SNMP v1: Set SNMP trap support to SNMP version 1.

SNMP v2c: Set SNMP trap support to SNMP version 2c (default).

SNMP v3: Set SNMP trap support to SNMP version 3.

Trap Community

Sets the community access string when sending an SNMP trap packet. The allowed string length is **0** to **255**, and the allowed content is ASCII characters from 33 to 126. The default is **'public'**.

Trap Destination Address

Sets the SNMP trap destination address. Enter a valid IP address in dotted decimal notation (x.y.z.w), with these restrictions:

- 1) x, y, z, and w must be decimal numbers from **0-255**,
- 2) x must not be 0 unless x, y, and w are also **0**,
- 3) x must not be **127**, and
- 4) x must not be greater than **223**.

Trap Destination IPv6 Address

Sets the trap destination IPv6 address of this LIB-44xx. 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, in the address '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 only appear once. It can also represent a legally valid IPv4 address (e.g., **::192.1.2.34**).

Trap Authentication Failure

Sets that the SNMP entity is permitted to generate authentication failure traps. The valid selections are:

Enabled: Enable SNMP trap authentication failure (default).

Disabled: Disable SNMP trap authentication failure.

Trap Link-up and Link-down

Sets the SNMP trap link-up and link-down mode of operation. The valid selections are:

Enabled: Enable SNMP trap link-up and link-down mode operation (default).

Disabled: Disable SNMP trap link-up and link-down mode operation.

Trap Inform Mode

Sets the SNMP trap inform mode operation. The valid selections are:

Enabled: Enable SNMP trap inform mode operation.

Disabled: Disable SNMP trap inform mode operation.

Trap Inform Timeout (seconds)

Sets the SNMP trap inform timeout. The valid range is **0** to **2147**. The default is **1**.

Trap Inform Retry Times

Sets the SNMP trap inform retry times. The valid range is **0** to **255**. The default is **5**.

Trap Probe Security Engine ID

Sets the SNMP V3 trap probe security engine ID mode of operation. The valid values are:

Enabled: Enable SNMP trap probe security engine ID mode of operation (default).

Disabled: Disable SNMP trap probe security engine ID mode of operation.

This field displays only if the 'Trap Version' parameter is set to **SNMP v3** (see the 'Trap Version' description above)

Trap Security Engine ID

Sets the SNMP trap security engine ID. SNMPv3 sends traps and informs using USM for authentication and privacy. A unique engine ID for these traps and informs is needed. When "Trap Probe Security Engine ID" is enabled, the ID will be probed automatically. Otherwise, the ID specified in this field is used. The string must contain an even number (in hexadecimal format) with between 10 and 64 characters, but all-zeros and all-'F's are not allowed. The message "Probe Fail" displays if the information could not be read.

This field displays only if the 'Trap Version' parameter is set to **SNMP v3** (see the 'Trap Version' description above)

Trap Security Name

Sets the SNMP trap security name. SNMPv3 traps and informs using USM for authentication and privacy.

A unique security name is needed when traps and informs are enabled. The default is **'None'**.

This field displays only if the 'Trap Version' parameter is set to **SNMP v3** (see the 'Trap Version' description above)

Buttons

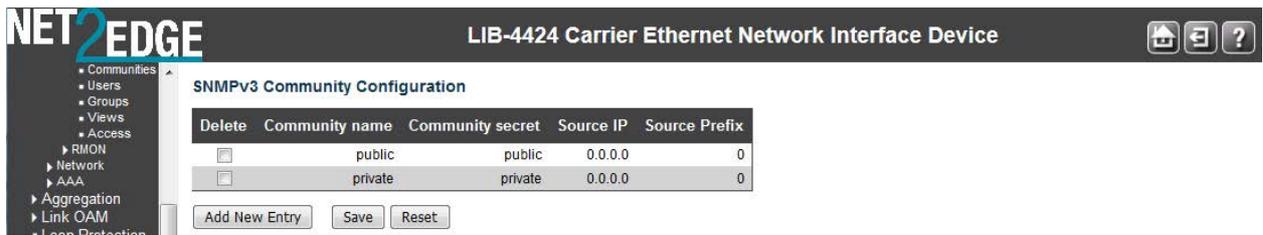
Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

SNMPv3 Community Configuration

Configure SNMPv3 community table on this page from the **Configuration > Security > Switch > SNMP > Communities** menu path. The entry index key is **Community**. SNMP V1 and V2c use a community string match for authentication. SNMP V3 uses a username match for authentication, or authentication based on the HMAC-MD5 or HMAC-SHA algorithms.

From the default page, click the **Add New Entry** button to display the new entry edit fields.



Delete

Check to delete the entry. It will be deleted during the next save.

Community

Indicates / sets the community access string to permit access to SNMPv3 agent. The allowed string length is 1 to 32 characters, and the allowed content is ASCII characters from 33 to 126 (e.g., space character not allowed). The community string will be treated as the security name and map an SNMPv1 or SNMPv2c community string.

Source IP

Indicates / sets the SNMP access source address. A particular range of source addresses can be used to restrict the source subnet when combined with source mask (e.g., **192.168.1.30**).

Source Mask

Indicates / sets the SNMP access source address mask (e.g., **255.255.255.0**). Enter a valid IP mask of a dotted decimal string ('x.y.z.w'), where:

- 1) x, y, z, and w are decimal numbers from **0-255**, and
- 2) when converted to a 32-bit binary string and read from left to right, all bits following the first zero must also be zero.

Buttons

Add New Entry: Click to add a new SNMP community entry. Enter the Community, Source IP, and Source Mask as described above.

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

SNMPv3 User Configuration

Configure SNMPv3 user table on this page from the **Configuration - Security - Switch - SNMP > User** menu path. The entry index keys are **Engine ID** and **User Name**.

The USM is supported per standard with a variety of user access levels and privacy protocols. SNMP v3 configuration involves setting up SNMP v3 Users, Groups, and Views. SNMP Users have a specified username, authentication password, privacy password, (if required) and authentication and privacy protocol assigned. The authentication protocol options are none, MD5, or SHA. The privacy algorithm options are none, AES, or DES. When a new User is created, it is associated with an SNMP group.

From the default page, click the **Add New Entry** button to display the new entry edit fields.

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 equals 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 33 to 126. No spaces can be entered.

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.

Note: if **Security Level** is set to **NoAuth, NoPriv**, then the remaining fields do not require an entry or selection.

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 that you 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 characters. 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.

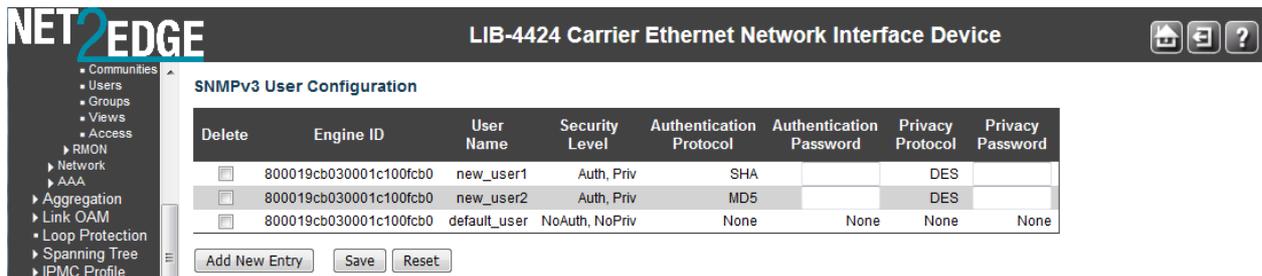
Privacy Password

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. No space characters are allowed.

Buttons

- Add New Entry:** Click to add a new user entry. Enter the Engine ID, User Name, Security Level, and Auth / Privacy entries as described above.
- Save:** Click to save changes.
- Reset:** Click to undo any changes made locally and revert to previously saved values.

The screen sample below shows two new SNMP v3 users added to the table.



SNMPv3 Group Configuration

Configure SNMPv3 group table on this page from the **Configuration > Security > Switch > SNMP > Group** menu path. The entry index keys are **Security Model** and **Security Name**.

SNMP v3 configuration involves setting up SNMP v3 Users, Groups, and Views. SNMP Groups are basically access control policies to which users can be added. Each SNMP Group is configured with a security model, and is associated with an SNMP security name. These parameters specify the type of authentication and privacy a user within the SNMP group will use, and also which objects in the MIB the User can access. Each SNMP Group name and security level pair must be unique within the device.

From the default page, click the **Add New Entry** button to display the new entry edit fields.

Delete	Security Model	Security Name	Group Name
<input type="checkbox"/>	v1	public	default_ro_group
<input type="checkbox"/>	v1	private	default_rw_group
<input type="checkbox"/>	v2c	public	default_ro_group
<input type="checkbox"/>	v2c	private	default_rw_group
<input type="checkbox"/>	usm	new_user1	default_rw_group

Delete

Check to delete the entry. It will be deleted during the next save.

Security Model

Indicates / sets the security model that this entry should belong to. Possible security models are:

v1: Reserved for SNMPv1.

v2c: Reserved for SNMPv2c.

usm: User-based Security Model (USM).

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 (e.g., no space characters).

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 (e.g., no space characters).

Buttons

Add New Entry: Click to add a new group entry to the table. Enter the Security Model, Security Name, and associated Group Name as discussed above.

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

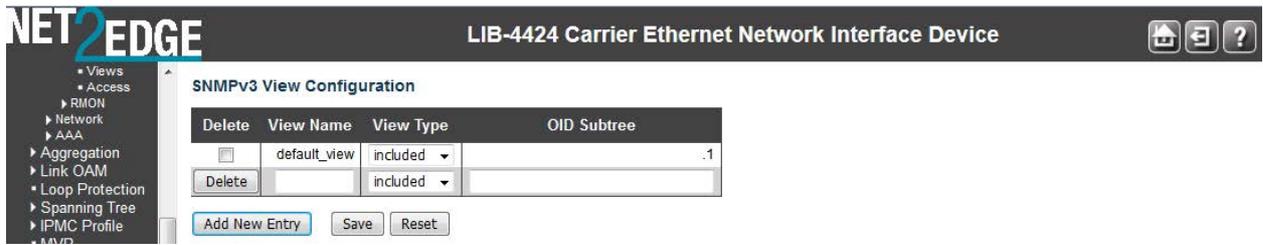
SNMPv3 Views Configuration

Configure SNMPv3 view table on this page from the **Configuration > Security > Switch > SNMP > Group** menu path. The entry index keys are View Name and OID Subtree.

SNMP v3 configuration involves setting up SNMP v3 Users, Groups, and Views. SNMP MIB Views are defined lists of objects within a MIB that can be used to control which parts of a MIB can be accessed by Users belonging to the SNMP Group associated with that particular View. Objects in the View may be from anywhere in the MIB, and are not required to be in the same MIB sub-tree.

When you have defined your Views, you must configure for your SNMP Groups the type of access Users will have to those Views.

From the default page, click the **Add New Entry** button to display the new entry edit fields.



Delete

Check the checkbox 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 its OID subtree should overstep the 'excluded' view entry.

OID Subtree

The OID defining the root of the subtree to add to the named view. The format of the OID Subtree is .OID1.OID2.OID3... The allowed OID length is **.1** to **.128**. The valid string content is a digital number or asterisk (*).

Buttons

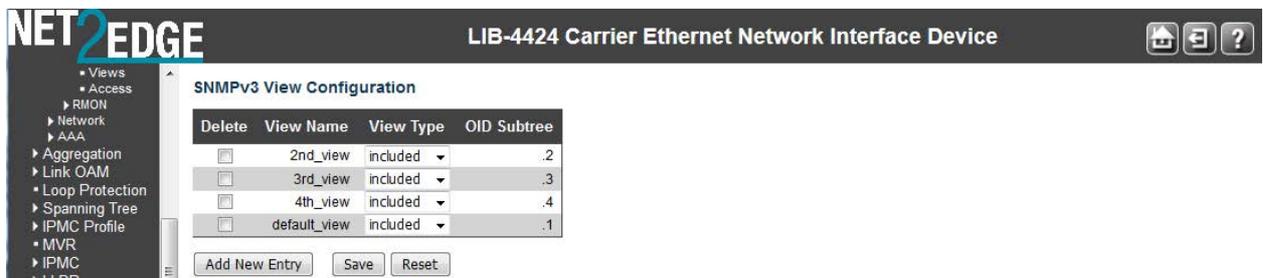
Add new view: Click to add a new View entry to the table. Enter the View Name, View Type, and OID Subtree as described above.

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Example

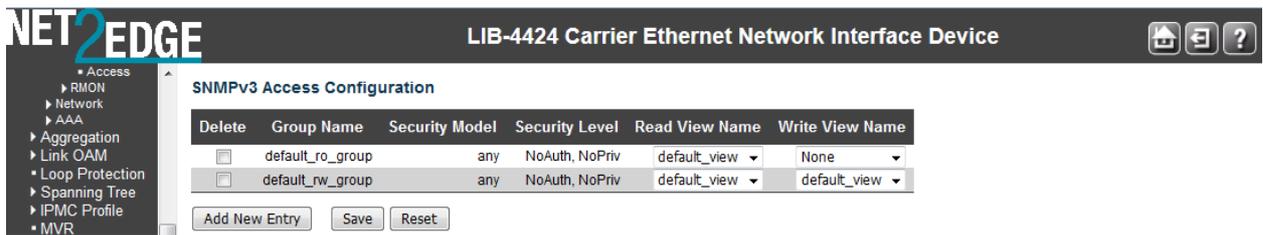
The screen below shows a second, third, and fourth SNMP v3 View added to the table.



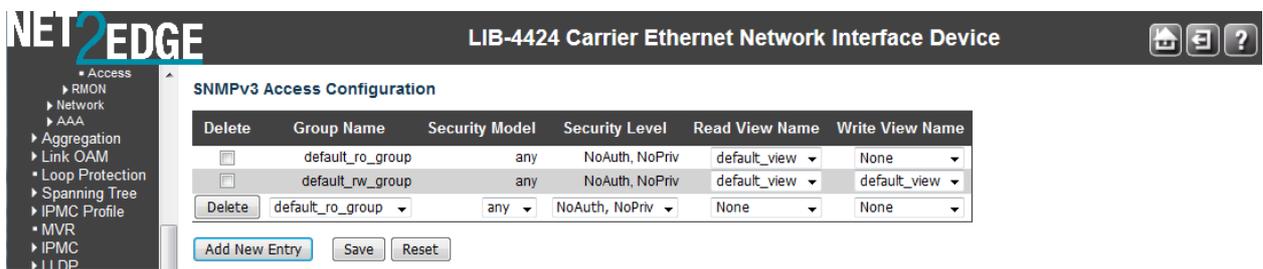
SNMPv3 Access Configuration

Configure the SNMPv3 access table on this page from the **Configuration > Security > Switch > SNMP > Access** menu path. The entry index keys are **Group Name**, **Security Model** and **Security**

Level. The default table displays two Groups: **default_ro_group** and **default_rw_group**. You can edit the **Read View Name** and the **Write View Name** for each of the two default entries in the initial table.



When you click the **Add New Entry** button, the SNMPv3 access table displays with fields for the new access group.



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 belongs to. Valid string lengths are **1** to **32**, and the allowed content is ASCII characters from 33 to 126. This dropdown lets you select an existing Group Name.

Security Model

Indicates the security model that this entry should belong to. Valid security models are:

any: Any SNMP security model accepted (v1, v2c, or usm).

v1: Reserved for SNMPv1.

v2c: Reserved for SNMPv2c.

usm: User-based Security Model (USM).

Security Level

Indicates the security model that this entry should belong to. Valid security models are:

NoAuth, NoPriv: No authentication and no privacy.

Auth, NoPriv: Authentication but 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.

The dropdown lets you select an existing Read View Name or **'None'**.

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.

The dropdown lets you select an existing Write View Name or '**None**'.

Buttons

Add New Entry: Click to add a new access entry to the table. Add the Group Name, Security Model, Security Level, Read View Name, and Write View Name as described above.

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Configuration > Security > Switch > RMON

Configure RMON Statistics, History, Alarms, and Events from the **Configuration > Security > Switch > RMON** menu path. The LIB-44xx RMON (Remote Network Monitoring) function supports the monitoring and protocol analysis of a LAN per [IETF RFC 1271](#). A part of SNMP, RMON is a network management protocol that gathers remote network information. See also "[RFC 1757 - Remote Network Monitoring Management Information Base](#)".

RMON collects nine kinds of information, including packets sent, bytes sent, packets dropped, statistics by host, by conversations between two sets of addresses, and certain kinds of events that have occurred. A network administrator can find out how much bandwidth or traffic each user is imposing on the network and what Web sites are being accessed. Alarms can be set to alert you of impending problems. The RMON Statistics, History, Alarm, and event data displays at **Monitor > Security > Switch > RMON > Event**.

RMON > Statistics

The **Configuration > Security > Switch > RMON > Statistics** menu path displays the **RMON Statistics Configuration** page. Configure RMON Statistics table on this page. The entry index key is **ID**.

The initial table displays no entries.

When you click the **Add New Entry** button, the RMON access table displays with fields for the new access group. You can edit the ID and the Data Source port suffix field in the table.

The RMON Statistics table parameters are explained below:

Delete

Click to delete the entry. It will be deleted during the next save.

ID

Enter the index for this entry. The valid range is **1** to **65535**.

Data Source

Indicates the port ID which you want to be monitored (e.g., .1.3.6.1.2.1.2.2.1.1.**0**). The valid range is **1-1013**.

Buttons

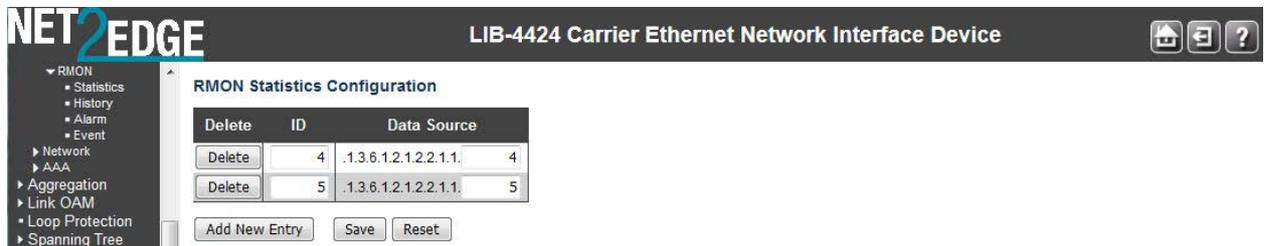
Add New Entry: Click to add a new community entry to the table. Add the ID and Data Source entries as described above.

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Example

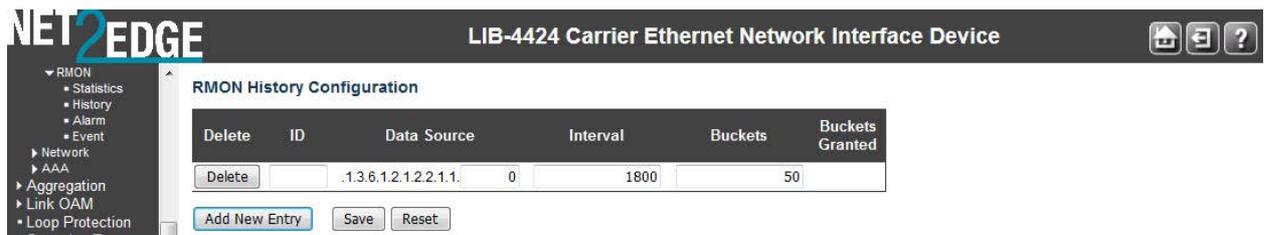
The example below shows two valid configured RMON statistics table entries.



RMON > History

The **Configuration > Security > Switch > SNMP > RMON > History** menu path displays the RMON History Configuration table. Configure RMON History table on this page. The entry index key is **ID**.

From the default screen, click the **Add New Entry** button to display the new entry fields.



The RMON History table parameters are explained below:

Delete

Check to delete the entry. It will be deleted during the next save.

ID

Indicates the index of the entry. The valid range is **1** to **65535**.

Data Source

Indicates the port ID which you want to be monitored with RMON.

Interval

Indicates the interval in seconds for sampling the history statistics data. The valid range is **1** to **3600**. The default value is **1800** seconds.

Buckets

Indicates the maximum data entries associated this History control entry stored in RMON. The valid range is **1** to **3600**. The default value is **50**. This is the RMON “buckets requested” value - the requested number of discrete time intervals over which data is to be saved in the part of the media-specific table associated with this historyControlEntry. When this object is created or modified, the probe should set historyControlBucketsGranted as closely to this object as is possible for the particular probe implementation and available resources. The default is **50**.

Buckets Granted

The number of data saved in the RMON. The number of discrete sampling intervals over which data will be saved in the part of the media-specific table associated with this *historyControlEntry*.

See the RMON RFC ([IETF RFC 2819](http://www.ietf.org/rfc/rfc2819.txt)) for details on the particular probe implementation and available resources.

Buttons

Add New Entry: Click to add a new community entry to the table. Enter the ID, Data Source, Interval, Buckets, and Buckets Granted parameters.

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Example

The example below shows two valid, saved RMON statistics table entries.

The screenshot shows the NET2EDGE web interface for a LIB-4424 Carrier Ethernet Network Interface Device. The left sidebar contains a navigation menu with options like RMON, Network, AAA, Aggregation, Link OAM, Loop Protection, Spanning Tree, and IPMC Profile. The main content area is titled "RMON History Configuration" and displays a table with the following data:

Delete	ID	Data Source	Interval	Buckets	Buckets Granted
Delete	1	.1.3.6.1.2.1.2.2.1.1.	1	900	50
Delete	2	.1.3.6.1.2.1.2.2.1.1.	2	500	50

Below the table are buttons for "Add New Entry", "Save", and "Reset".

RMON > Alarm

The **Configuration > Security > Switch > SNMP > RMON > Alarm** menu path displays the RMON Alarm Configuration table. Configure RMON Alarm table on this page. The entry index key is **ID**.

When you click the **Add New Entry** button from the default table page, the table displays with entry fields.

The screenshot shows the NET2EDGE web interface for a LIB-4424 Carrier Ethernet Network Interface Device. The left sidebar contains a navigation menu with options like RMON, Network, AAA, Aggregation, Link OAM, Loop Protection, Spanning Tree, and IPMC Profile. The main content area is titled "RMON Alarm Configuration" and displays a table with the following data:

Delete	ID	Interval	Variable	Sample Type	Value	Startup Alarm	Rising Threshold	Rising Index	Falling Threshold	Falling Index
Delete		30	.1.3.6.1.2.1.2.2.1.1. 0.0	Delta	0	RisingOrFalling	0	0	0	0

Below the table are buttons for "Add New Entry", "Save", and "Reset".

The RMON Alarm Configuration table parameters are explained below:

Delete

Click to delete the entry on this line. If not previously saved, it will be deleted immediately; otherwise, it will be deleted during the next save.

ID

Indicates / set the port index of the entry. The valid range is **1** to **65535**.

Interval

Indicates the interval in seconds for sampling and comparing the rising and falling threshold. The valid range is **1** to **2³¹-1**.

Variable

Enter a variable value in the format xxx.yyy, where xxx is 10-21, and yyy is 1-65,535.

Indicates the particular variable to be sampled. The valid variables are:

InOctets: The total number of octets received on the interface, including framing characters.

InUcastPkts: The number of uni-cast packets delivered to a higher-layer protocol.

InNUcastPkts: The number of broadcast and multi-cast packets delivered to a higher-layer protocol.

InDiscards: The number of inbound packets that are discarded even the packets are normal.

InErrors: The number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol.

InUnknownProtos: the number of the inbound packets that were discarded because of the unknown or un-support protocol.

OutOctets: The number of octets transmitted out of the interface, including framing characters.

OutUcastPkts: The number of unicast packets that request to transmit.

OutNUcastPkts: The number of broadcast and multi-cast packets that request to transmit.

OutDiscards: The number of outbound packets that are discarded in the event the packet is normal.

OutErrors: The number of outbound packets that could not be transmitted because of errors.

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. The valid 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 - 2147483647). The Rising Threshold must be larger than the Falling Threshold.

Rising Index

Rising event index (1 - 65535). The Rising Threshold must be larger than the Falling Threshold.

Falling Threshold

Falling threshold value (-2147483648 - 2147483647). The Falling Threshold must be smaller than the Rising Threshold.

Falling Index

Falling event index (1 - 65535). The Falling Index must be smaller than the Rising Index.

Buttons

Add new entry: Click to add a new community entry.

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Example

A table with two new RMON Alarm Configuration entries is shown below:

The screenshot shows the NET2EDGE interface for a LIB-4424 Carrier Ethernet Network Interface Device. The main content area displays the 'RMON Alarm Configuration' table. The table has columns for Delete, ID, Interval, Variable, Sample Type, Value, Startup Alarm, Rising Threshold, Rising Index, Falling Threshold, and Falling Index. Two entries are shown:

Delete	ID	Interval	Variable	Sample Type	Value	Startup Alarm	Rising Threshold	Rising Index	Falling Threshold	Falling Index
Delete	1	30	.1.3.6.1.2.1.2.2.1.11.111	Delta	0	RisingOrFalling	0	0	0	0
Delete	2	30	.1.3.6.1.2.1.2.2.1.22.222	Delta	0	RisingOrFalling	0	0	0	0

Below the table are buttons for 'Add New Entry', 'Save', and 'Reset'.

RMON > Event

The **Configuration > Security > Switch > SNMP > RMON > Event** menu path displays the RMON Alarm Configuration table. Configure RMON Event table on this page. The entry index key is **ID**.

When you click the **Add New Entry** button from the default table page, the table displays with entry fields.

The screenshot shows the NET2EDGE interface for a LIB-4424 Carrier Ethernet Network Interface Device. The main content area displays the 'RMON Event Configuration' table. The table has columns for Delete, ID, Desc, Type, and Event Last Time. One entry is shown:

Delete	ID	Desc	Type	Event Last Time
Delete			none	0

Below the table are buttons for 'Add New Entry', 'Save', and 'Reset'.

The RMON Event table parameters are explained below:

Delete

Click to delete the entry. If not previously saved, it will be deleted immediately; otherwise it will be deleted during the next save.

ID

Enter the index of the RMON event. The valid range is **1** to **65535**. Each ID entry must be unique.

Desc

Indicates this event, the string length is **0** to **127**. The default is a **null** string.

Type

Indicates the notification of the event, the valid types are:

none: No logging action is performed.

log: A syslog entry is added.

snmptrap: A SNMP trap event is sent.

logandtrap: A syslog entry is logged and an SNMP trap event is sent.

Community

Specify the community when a trap is sent, the string length is **0** to **127 characters**. The default is "**public**".

Event Last Time

Indicates the value of sysUpTime at the time this event entry last generated an event (e.g., 33554560 or 33 days, 55 hours, 45 minutes, and 50 seconds).

Buttons

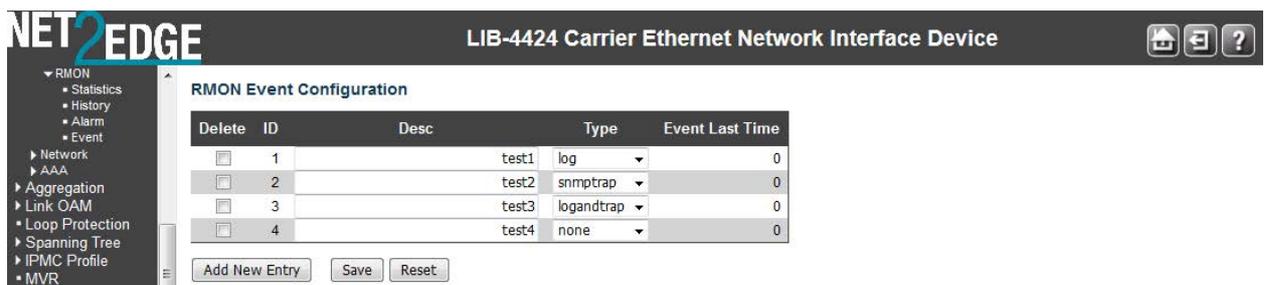
Add New Entry: Click to add a new community entry to the table. Enter the ID, Desc, Type, Community, and Event Last Time values as described above.

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Example

An RMON Event Configuration table with four entries is shown below (IDs 1-4).



Note that you can monitor the related RMON Statistics, History, Alarm, and Event data from the **Monitor > Security > Switch > RMON > Event** menu path.

Port Security Limit Control Configuration

Configuration > Security > Network > Limit Control

The **Configuration > Security > Network > Limit Control** menu path lets you to configure the Port Security Limit Control system-level and port-level settings.

Limit Control allows for limiting the number of users on a given port. A user is identified by a MAC address and VLAN ID. If Limit Control 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. The action can be one of the four different actions described below:

The Limit Control module utilizes a lower-layer module (the Port Security module) which manages MAC addresses learned on the port. The Limit Control configuration consists of two sections, a system-wide and a port-wide configuration table.

The screenshot shows the NET2EDGE web interface for a LIB-4424 Carrier Ethernet Network Interface Device. The main content area is titled "Port Security Limit Control Configuration" and includes a "Refresh" button. The configuration is split into two sections:

System Configuration

- Mode:** Disabled (dropdown menu)
- Aging Enabled:**
- Aging Period:** 3600 seconds

Port Configuration

Port	Mode	Limit	Action	State	Re-open
*	<>	4	<>		
1	Disabled	4	None	Disabled	Reopen
2	Disabled	4	None	Disabled	Reopen
3	Disabled	4	None	Disabled	Reopen
4	Disabled	4	None	Disabled	Reopen
5	Disabled	4	None	Disabled	Reopen

System Configuration

Mode

Indicates if Limit Control is globally enabled or disabled on the LIB-44xx. If globally disabled, other modules may still use the underlying functionality, but limit checks and corresponding actions are disabled.

Aging Enabled

If checked, secured MAC addresses are subject to aging as discussed under 'Aging Period' below: To keep the MAC table updated, an aging scan is conducted to remove entries that were not recently accessed. This ensures that stations moved to new locations are not permanently prevented from receiving frames in their new location. It also frees up MAC table entries occupied by obsolete stations to make room for new stations. The IEEE 802.1d recommends 300 seconds per entry.

Aging Period

If checked, then the aging period is controlled with this input. If other modules are using the underlying port security for securing MAC addresses, they may have other requirements to the aging period. The underlying port security will use the shorter requested aging period of all modules that use the functionality.

The Aging Period can be set to a number between **10** and **10,000,000** seconds. The IEEE 802.1d recommends **300** seconds per entry.

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 Limit Control 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.

Port Configuration

At Port Security Limit Control Configuration, the Port Configuration table has one row for each LIB-44xx port and a number of columns, which are explained below:

Port

The port number to which the configuration below applies. The * in the Port column acts as a 'wild card' character which causes the selections in this row to be applied to all other rows in the table for which this selection is valid. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Mode

Controls whether Limit Control is enabled on this port. Both this and the 'Global Mode' must be set to Enabled for Limit Control to be in effect. Notice that other modules may still use the underlying port security features without enabling Limit Control on a given port.

Limit

The maximum number of MAC addresses that can be secured on this port. Enter a number from **1** - **1023**.

If this limit is exceeded, the corresponding action is taken. The LIB-44xx 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.

Action

If the 'Limit' defined above is reached, the LIB-44xx can take one of the following actions:

None: Do not allow more than 'Limit' MAC addresses on the port, but take no further action.

Trap: If 'Limit' + 1 MAC addresses is seen on the port, send an SNMP trap. If Aging is disabled, only one SNMP trap will be sent, but with Aging enabled, new SNMP traps will be sent every time the limit gets exceeded.

Shutdown: If 'Limit' + 1 MAC addresses is seen on the port, shut down the port. This implies that all secured MAC addresses will be removed from the port, and no new address will be learned. Even if the link is physically disconnected and reconnected on the port (by disconnecting the cable), the port will remain shut down. There are three ways to re-open the shutdown port:

- 1) Boot the LIB-44xx.
- 2) Disable and re-enable Limit Control on the port or the LIB-44xx, or
- 3) Click the Reopen button.

Trap & Shutdown: If 'Limit' + 1 MAC addresses is seen on the port, both the "Trap" and the "Shutdown" actions described above will be taken.

State

This column shows the current state of the port as seen from the Limit Control's point of view. The state takes one of four values:

Disabled: Limit Control is either globally disabled or disabled on the port.

Ready: The limit is not yet reached. This can be shown for all actions.

Limit Reached: Indicates that the limit is reached on this port. This state can only be shown if 'Action' is set to **None** or **Trap**.

Shutdown: Indicates that the port is shut down by the Limit Control module. This state can only be shown if 'Action' is set to **Shutdown** or to **Trap & Shutdown**.

Re-open Button

If a port is shutdown by this module, you may reopen it by clicking this button, which will only be enabled if this is the case. For other methods, refer to the **Shutdown** description in the Action section above.

Note: clicking the **Reopen** button causes the page to be refreshed, so non-committed (unsaved) changes will be lost.

Buttons

Refresh: Click to refresh the page. Note that non-committed (unsaved) changes will be lost.

Reset: Click to undo any changes made locally and revert to previously saved values.

Save: Click to save changes.

Example

The screen below shows Port Security Limit Control Configuration changes after doing a **Save**.

NET2EDGE LIB-4424 Carrier Ethernet Network Interface Device

Port Security Limit Control Configuration

System Configuration

Mode: Enabled

Aging Enabled:

Aging Period: 600 seconds

Port Configuration

Port	Mode	Limit	Action	State	Re-open
*	<>	4	<>		
1	Enabled	4	Trap	Ready	Reopen
2	Enabled	4	Shutdown	Ready	Reopen
3	Enabled	4	Trap & Shutdown	Ready	Reopen
4	Enabled	4	Trap	Ready	Reopen
5	Enabled	4	Shutdown	Ready	Reopen

NAS (Network Access Server) Configuration

Configuration > Security > Network > NAS

The **Configuration > Security > Network > NAS** menu path lets you 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, the backend servers, determine whether the user is allowed access to the network.

These backend (RADIUS) servers are configured from the **Configuration > Security > AAA** menu path.

The IEEE802.1X standard defines port-based operation, but non-standard variants overcome security limitations as explained 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. A device 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.

IEEE 802.1X Port-based Network Access Control provides a standard for authenticating and authorizing devices attached to a LAN port. Generally, IEEE 802.1X is port-based; however, the LIB-44xx also supports MAC-based network access control.

The NAS configuration consists of two sections, for system-wide and port-wide NAS configuration.

NET2EDGE LIB-4424 Carrier Ethernet Network Interface Device

Network Access Server Configuration [Refresh]

System Configuration

Mode	Disabled
Reauthentication Enabled	<input type="checkbox"/>
Reauthentication Period	3600 seconds
EAPOL Timeout	30 seconds
Aging Period	300 seconds
Hold Time	10 seconds
RADIUS-Assigned QoS Enabled	<input type="checkbox"/>
RADIUS-Assigned VLAN Enabled	<input type="checkbox"/>
Guest VLAN Enabled	<input type="checkbox"/>
Guest VLAN ID	1
Max. Reauth. Count	2
Allow Guest VLAN if EAPOL Seen	<input type="checkbox"/>

Port Configuration

Port	Admin State	RADIUS-Assigned QoS Enabled	RADIUS-Assigned VLAN Enabled	Guest VLAN Enabled	Port State	Restart
*	<>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Globally Disabled	
1	Force Authorized	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Globally Disabled	Reauthenticate Reinitialize
2	Force Authorized	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Globally Disabled	Reauthenticate Reinitialize
3	Force Authorized	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Globally Disabled	Reauthenticate Reinitialize

The NAS page parameters are explained below:

System Configuration

Mode

Indicates if NAS is globally enabled or disabled on the LIB-44xx. 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

Sets 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 on MAC-based ports.

Aging Period

This setting applies to the following modes (i.e., modes using the Port Security function) 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 determine 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 determine 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 **1** - **255**.

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 **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 table has one row for each LIB-44xx port and a number of columns, which are explained below:

Port

The port number for which the configuration below applies. The * in the Port column acts as a 'wild card' character which causes the selections in this row to be applied to all other rows in the table for which this selection is valid. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Admin State

If NAS is globally enabled, this selection controls the port's authentication mode. Note that the 802.1x Admin State must be set to **Force Authorize** for ports enabled for Spanning Tree. The Spanning Tree function is configured at **Configuration > Spanning Tree > CIST Ports > CIST Normal Port Configuration** in the "STP Enabled" column. The following modes are available:

Force Authorized: In this mode, the switch will send one EAPOL Success frame when the port link comes up, and any client on the port will be allowed network access without authentication.

Force Unauthorized: In this mode, the switch will send one EAPOL Failure frame when the port link comes up, and any client on the port will be disallowed network access.

Port-based 802.1X: In 802.1X, the user is called the supplicant, the switch is the authenticator, and the RADIUS server is the authentication server. The authenticator acts as the man-in-the-middle, forwarding requests and responses between the supplicant and the authentication server. Frames sent between the supplicant and the switch are special 802.1X frames, known as EAPOL (EAP Over LANs) frames. EAPOL frames encapsulate EAP PDUs ([RFC3748](#)). Frames sent between the switch and the RADIUS server are **RADIUS** packets. RADIUS packets also encapsulate EAP PDUs together with other attributes like the switch's IP address, name, and the supplicant's port number on the switch. EAP is very flexible, in that it allows for different authentication methods, like [MD5-Challenge](#), [PEAP](#), and [TLS](#). The important thing is that the authenticator (the switch) doesn't need to know which authentication method the supplicant and the authentication server are using, or how many information exchange frames are needed for a particular method. The switch simply encapsulates the EAP part of the frame into the relevant type (EAPOL or RADIUS) and forwards it.

When authentication is complete, the RADIUS server sends a special packet containing a success or failure indication. Besides forwarding this decision to the supplicant, the switch uses it to open up or block traffic on the switch port connected to the supplicant.

Note: Suppose two backend servers are enabled and that the server timeout is configured to X seconds (using the AAA configuration page), and suppose that the first server in the list is currently down (but not considered dead). Now, if the supplicant retransmits EAPOL Start frames at a rate faster than X seconds, then it will never get authenticated, because the switch will cancel on-going backend authentication server requests whenever it receives a new EAPOL Start frame from the supplicant. And since the server hasn't yet failed (because the X seconds haven't expired), the same server will be contacted upon the next backend authentication server request from the switch. This scenario will loop forever. Therefore, the server timeout should be smaller than the supplicant's EAPOL Start frame retransmission rate.

Single 802.1X: In port-based 802.1X authentication, once a supplicant is successfully authenticated on a port, the whole port is opened for network traffic. This allows other clients connected to the port (for instance through a hub) to piggy-back on the successfully authenticated client and get network access even though they really aren't authenticated. To overcome this security breach, use the Single 802.1X variant.

Single 802.1X is really not an IEEE standard, but features many of the same characteristics as does port-based 802.1X. In Single 802.1X, at most one supplicant can get authenticated on the port at a time. Normal EAPOL frames are used in the communication between the supplicant and the switch. If more than one supplicant is connected to a port, the one that comes first when the port's link comes up will be the first one considered. If that supplicant doesn't provide valid credentials within a certain amount of time, another supplicant will get a chance. Once a supplicant is successfully authenticated, only that supplicant will be allowed access. This is the most secure of all the supported modes. In this mode, the Port Security module is used to secure a supplicant's MAC address once successfully authenticated.

Multi 802.1X: In port-based 802.1X authentication, once a supplicant is successfully authenticated on a port, the whole port is opened for network traffic. This allows other clients connected to the port (for instance through a hub) to piggy-back on the successfully authenticated client and get network access even though they really aren't authenticated. To overcome this security breach, use the Multi 802.1X variant. Multi 802.1X is really not an IEEE standard, but features many of the same characteristics as does port-based 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" or "xx.xx.xx.xx.xx.xx" or "xxxxxxxxxx" (x is a hexadecimal digit). 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 port-based 802.1X is that several clients can be connected to the same port (e.g., through a 3rd party switch or a hub) and still require individual authentication, and that the clients don't need special supplicant software to authenticate. 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.

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:

- Port-based 802.1X
- Single 802.1X

RADIUS attributes used in identifying a QoS Class:

Refer to the IETF document for a description of the RADIUS attributes needed to successfully identify 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' - ', 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:

- Port-based 802.1X
- Single 802.1X

For troubleshooting VLAN assignments, use the **Monitor > VLANs > VLAN Membership** and the **Monitor > VLANs > VLAN Port** menu paths. These pages show which modules have (temporarily) overridden the current Port VLAN configuration.

RADIUS attributes used in identifying a VLAN ID:

IETF [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- 4094].

Guest VLAN Enabled

When Guest VLAN is both globally enabled and enabled (checked) for a given port, the switch considers moving the port into the Guest VLAN according to the rules outlined below:

This option is only available for EAPOL-based modes, i.e.:

- Port-based 802.1X
- Single 802.1X
- Multi 802.1X

For trouble-shooting VLAN assignments, use the **Monitor > VLANs > VLAN Membership** and the **Monitor > VLANs > VLAN Port** pages. These pages show which modules have (temporarily) overridden the current Port VLAN configuration.

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 the 'Max. Reauth. Count' and no EAPOL frames have been received in the meanwhile, the switch considers entering the Guest VLAN. The interval between transmissions of EAPOL Request Identity frames is configured with

'EAPOL Timeout. If Allow Guest VLAN if EAPOL Seen' 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 in the 'Restart' column 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.

This 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.



Buttons

Refresh: Click to refresh the page.

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Messages

Message: The 802.1X Admin State must be set to Authorized for ports that are enabled for LACP

Meaning: You made a change to an unsupported configuration.

Recovery:

1. Click the browser back button and re-configure the NAS settings, or
2. At the **Configuration > Spanning Tree > CIST Ports** menu path, disable

Message: The 802.1X Admin State must be set to Authorized for ports that are enabled for Spanning Tree

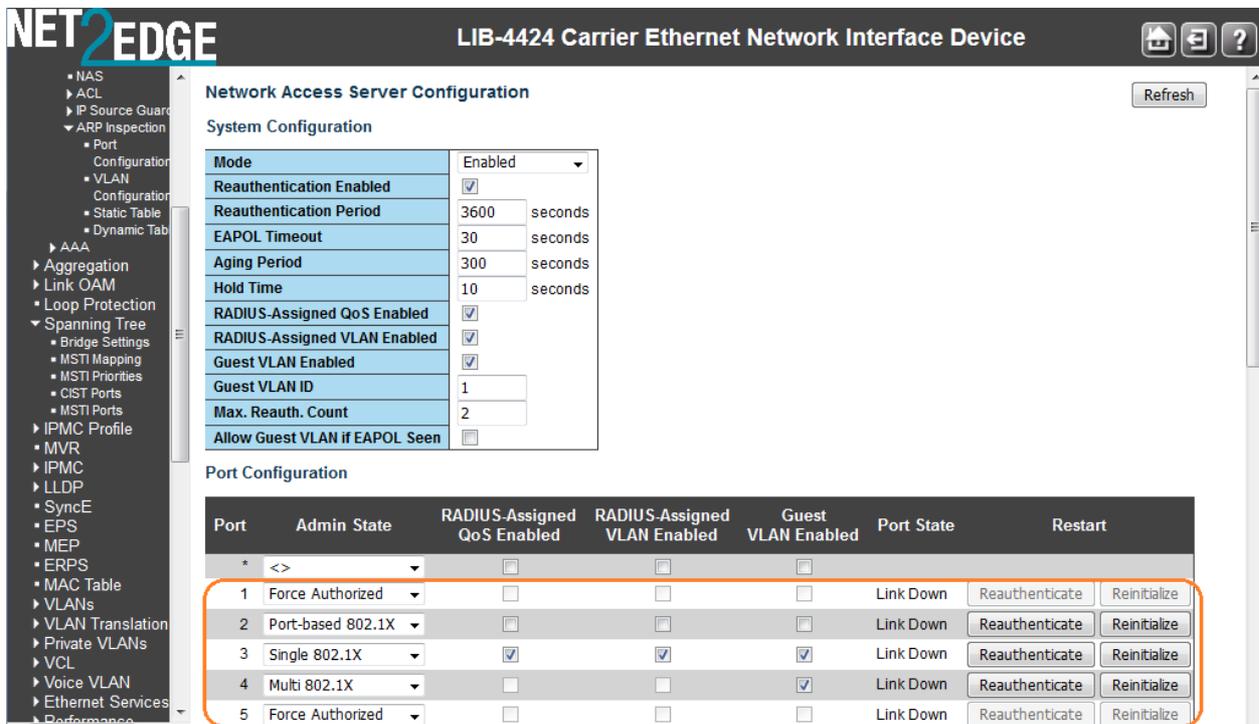
Meaning: You made a change to an unsupported configuration.

Recovery:

1. Click the browser back button and re-configure the NAS settings, or
2. At the **Configuration > Aggregation > LACP** menu path, disable LACP.

Example

In the sample **Configuration > Security > NAS** setup below, following a Save, Port 3 shows “Single 802.1X”, RADIUS-assigned QoS disabled / VLAN disabled, Guest VLAN enabled, Authorized port state, with “Reauthenticate” and “Reinitialize” restart enabled.



The 802.1X Admin State must be set to ‘Authorized’ for ports that are enabled for Spanning Tree. You can disable STP at the port level at **Configuration > Spanning Tree > CIST Port** menu path by unchecking the “STP Enabled” checkbox.

Note that the two buttons in the ‘Restart’ column are available for the Port 3 row. The **Reauthenticate** and **Reinitialize** buttons are only enabled when authentication is globally enabled and the port's Admin State is in an EAPOL-based or MAC-based mode.

ACL Ports Configuration

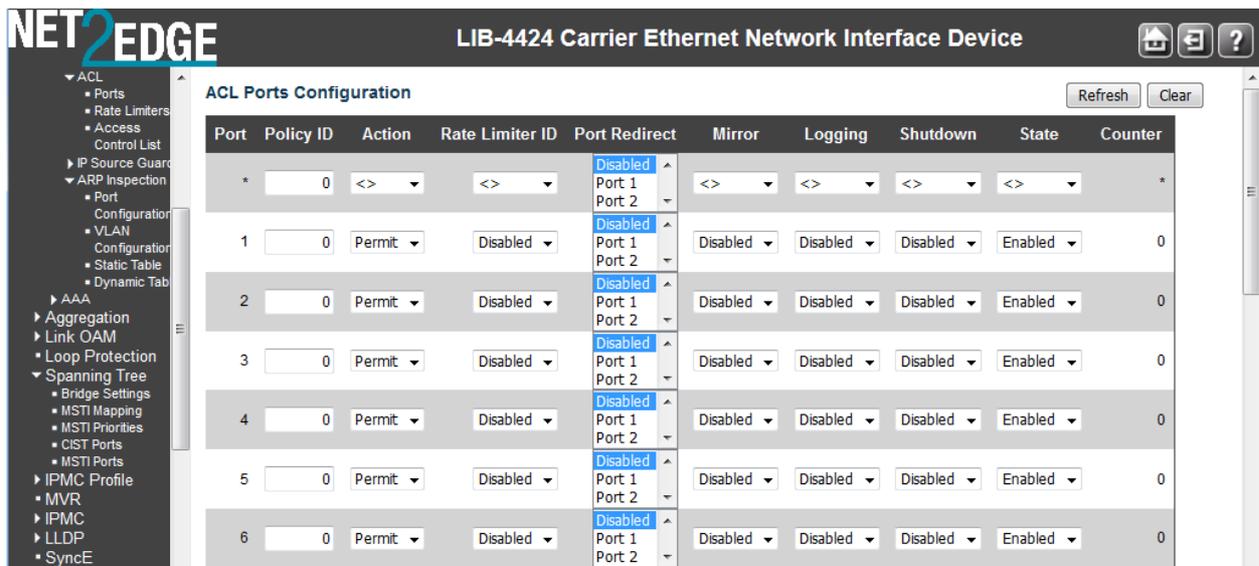
Configure the ACL parameters (ACE) of each LIB-44xx port from the **Configuration > Security > Network > ACL > Ports** menu path. These parameters will affect frames received on a port unless the frame matches a specific ACE.

Access Controls Lists

The LIB-44xx can ‘peek’ into the frames at line rate and is capable of deep packet inspection; this ability gives a wide range of access controls. The rules or the access control lists can look at any field in the Layer 2 to Layer 4 headers to make the decision of allowing, discarding, mirroring, logging or even shutdown the port that the frame came through.

The ACL rule created can be associated with any port as well when created as a policy. Apart from the ACL, there is a device level option to do storm prevention for the unicast, multicast and broadcast frames.

ACE (Access Control Entry) describes access permissions 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 application.



The **Configuration > Security > Network > ACL > Ports** page parameters are explained below:

Port

The logical port for the settings contained in the same row. The * in the **Port** column acts as a ‘wild card’ character which causes the selections in this row to be applied to all other rows in the table for which this selection is valid. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Policy ID

Select the policy to apply to this port. The allowed values are **1 - 8**. The default value is **1**.

Action

Select whether forwarding is permitted (**Permit**) or denied (**Deny**). The default value is **Permit**.

Note that **Action** can not be set to **Permit** with a 'Port Copy' setting of other than Disabled.

Rate Limiter ID

Select which rate limiter to apply on this port. The allowed values are **Disabled** or the values **1 - 16**. The default value is **Disabled**.

Port Redirect

Select which port frames are copied on. The allowed values are **Disabled** or a specific port number (i.e., **Port 1 - Port 4**).

The default value is **Disabled**. Note that 'Action' cannot be set to Permit with a Port Redirect setting of other than Disabled.

Logging

Specify the logging operation of this port. 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". Please note that 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".

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. This is a 'read only' field.

Buttons

Save: Click to save changes.

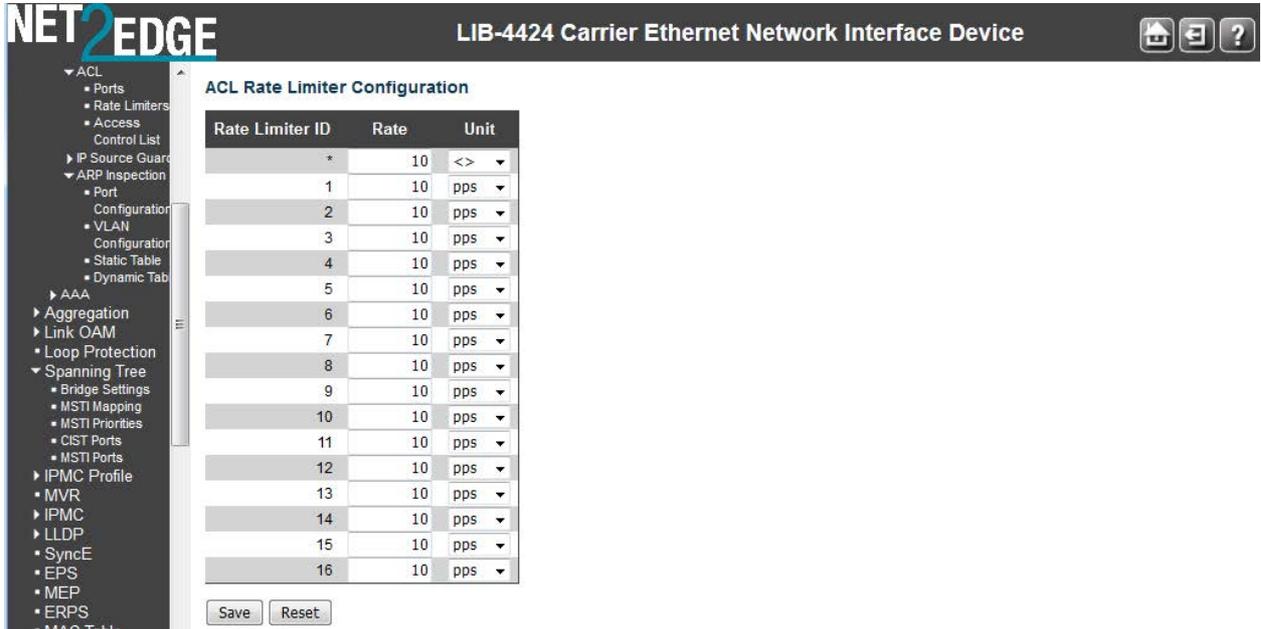
Reset: Click to undo any changes made locally and revert to previously saved values.

Refresh: Click to refresh the page; any changes made locally will be undone.

Clear: Click to clear the counters.

ACL Rate Limiter Configuration

Configure the rate limiters for the ACL of the LIB-44xx from the **Configuration > Security > Network > ACL > Rate Limiters** menu path.



Rate Limiter ID

The rate limiter ID for the settings contained in the same row. The * in the Rate Limiter ID column acts as a ‘wild card’ character which causes the selections in this row to be applied to all other rows in the table for which this selection is valid.

Rate

The valid values are: **0-131071** in pps (packets per second). The default value is **1**.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previous (unsaved) values.

Access Control List (ACL) Configuration

This page shows the Access Control List (ACL), which is made up of the ACEs defined on this LIB-44xx.

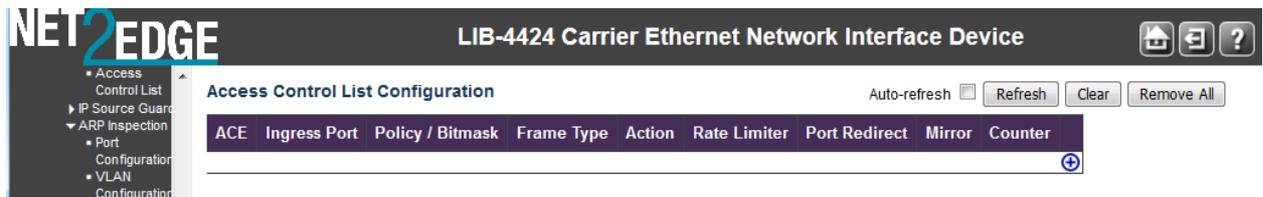
Each row describes the ACE that is defined. The maximum number of ACEs is **256** on each LIB-44xx.

Click on the lowest plus sign to add a new ACE to the list. The reserved ACEs used for internal protocol cannot be edited or deleted, the order sequence cannot be changed, and the priority is highest.

Each ACE (Access Control Entry) describes access permissions 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, varied parameter options that are available for individual application.

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.

A frame that hits this ACE matches the configuration that is defined here.



Ingress Port

Indicates the ingress port of the ACE. Valid 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.

Rate Limiter

Indicates the rate limiter number of the ACE. The valid range is **1** to **16**. When **Disabled** is displayed, rate limiter operation is disabled.

Port Redirect

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. Use this button initially (from the default screen) to create an initial ACE.

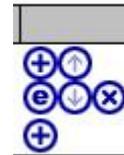
: Edits the ACE row.

: Moves the ACE up the list.

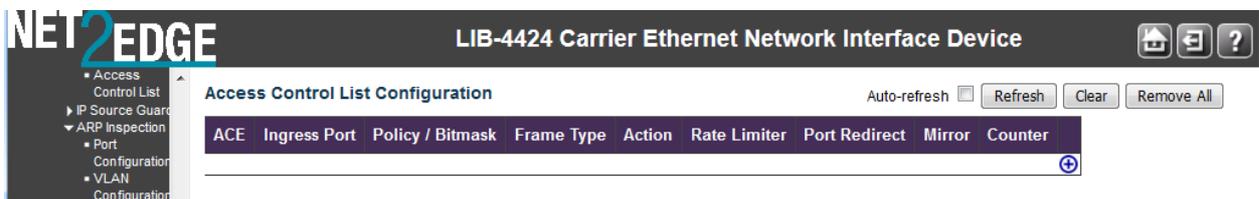
: Moves the ACE down the list.

: Deletes the ACE.

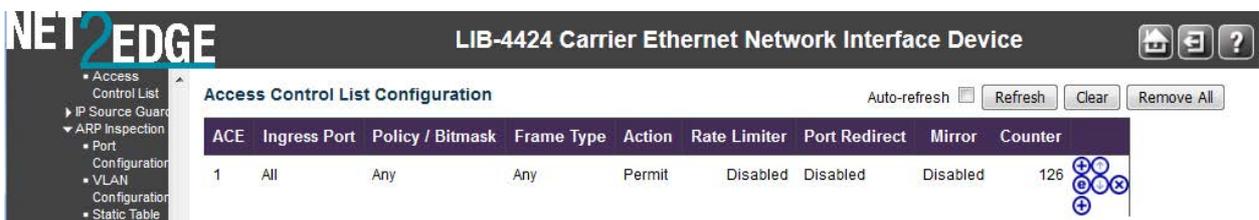
: The lowest plus sign adds a new entry at the bottom of the ACE listings.



These Modification buttons display in the far-right column of the Access Control List Configuration table.



When you click  to add the initial ACE row, the ACE row edit screen displays to let you enter the above parameters for an initial ACE entry.



Buttons

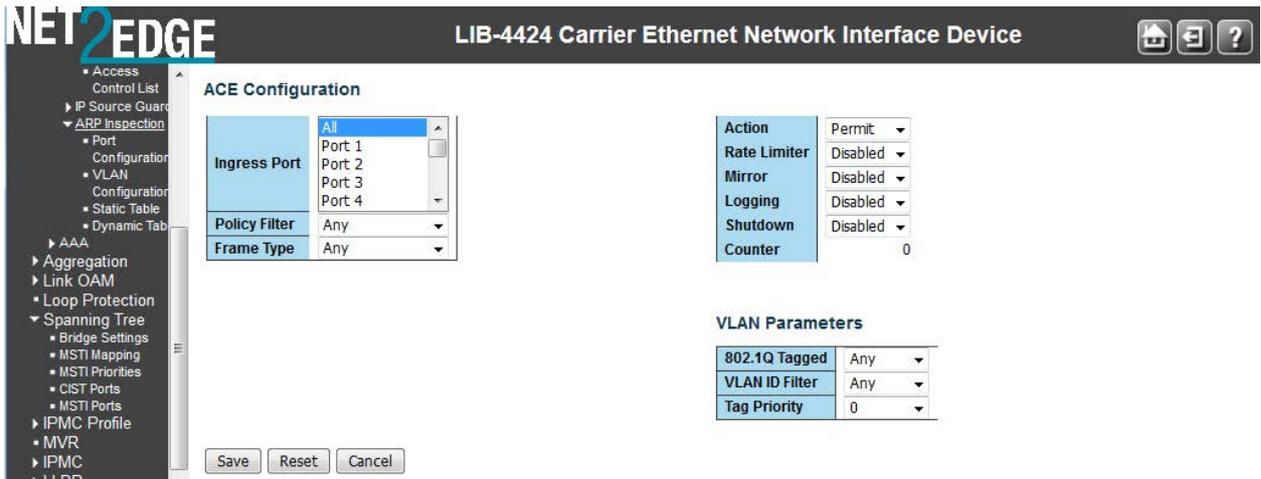
Auto-refresh: Check to refresh the page automatically. Automatic refresh occurs every three seconds.

Refresh: Click to refresh the page; any changes made locally will be undone.

Clear: Click to clear the Counter column in the table.

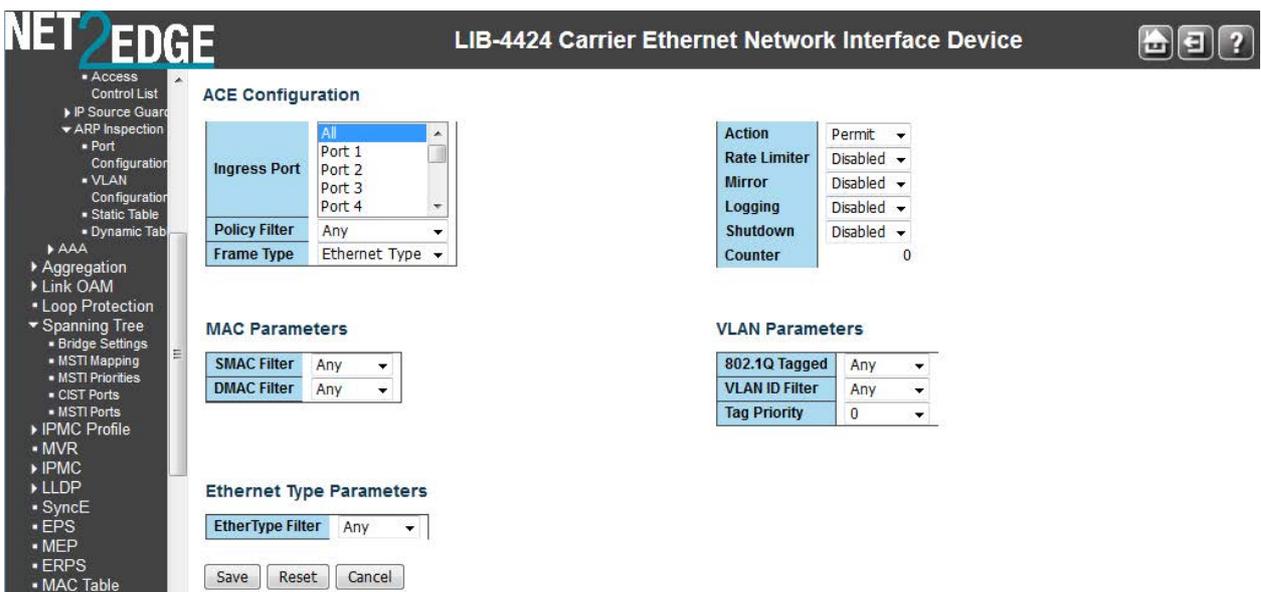
Remove All: Click to remove all existing ACE entries from the tables.

When you click the  button to edit an existing ACE row, the ACE row edit screen displays to let you edit the above parameters for an existing ACE entry.

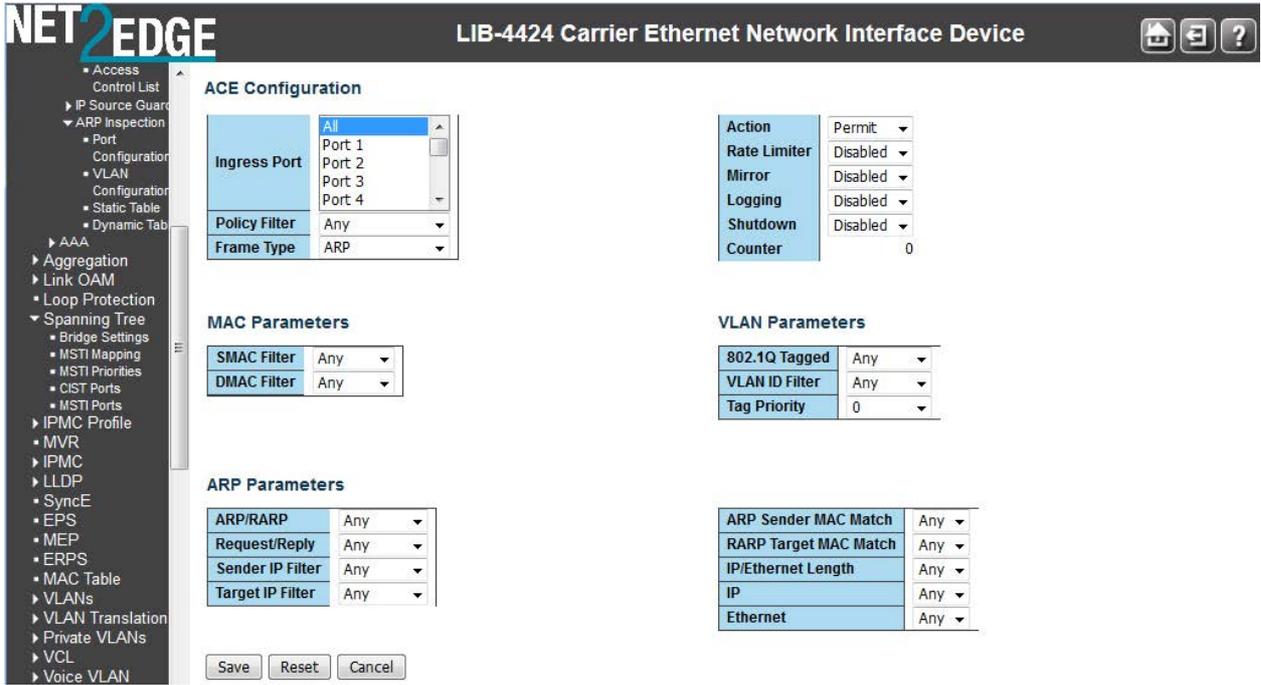


The parameters displayed depend on the config selections. For example, the config screen above displays if you select **Frame Type = Any**.

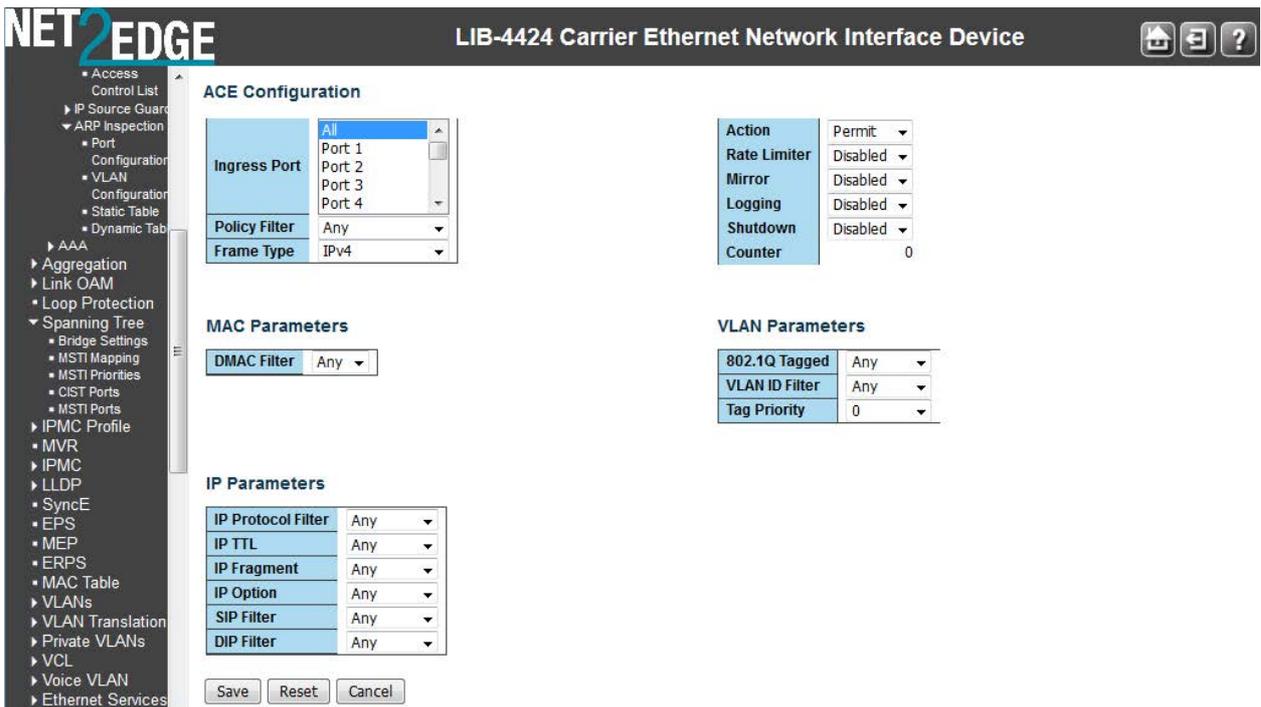
The configuration screen below displays if you select **Frame Type = Ethernet Type**.



The configuration screen below displays if you select **Frame Type = ARP**.



The configuration screen below displays if you select **Frame Type = IPv4**.



The various ACL parameters are explained below:

ACE Configuration Parameters

The ACE Configuration parameters let you configure the Access Control List Configuration table parameters as described earlier in this section

Ingress Port	All
Policy Filter	All
Frame Type	IPv4

Ingress Port	All
Policy Filter	Any
Frame Type	Any

Ingress Port	All
Policy Filter	Any
Frame Type	IPv4

MAC Parameters

SMAC Filter

Specify the source MAC filter for this ACE. (Only displayed when the frame type is Ethernet Type or ARP.)

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 an SMAC value displays.

MAC Parameters

SMAC Filter	Any
DMAC Filter	Any

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-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.

MAC Parameters

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

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.

MAC Parameters

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

VLAN Parameters

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 displays.

VLAN Parameters

VLAN ID Filter	Any
Tag Priority	0

VLAN ID

When "Specific" is selected for the VLAN ID filter, you can enter a specific VLAN ID number. The valid range is **1** to **4094**.

A frame that hits this ACE matches this VLAN ID value.

VLAN Parameters

VLAN ID Filter	Specific
VLAN ID	0
Tag Priority	0

Tag Priority

Specify the tag priority for this ACE. A frame that hits this ACE matches this tag priority. The valid number range is **0** to **7**. The value **Any** means that no tag priority is specified (tag priority is "don't-care".)

ARP Parameters

The ARP parameters can be configured when Frame Type "ARP" is selected.

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/RARP opcode set to ARP.

RARP: Frame must have ARP/RARP opcode set to RARP.

Other: Frame has unknown ARP/RARP Opcode flag.

ARP Parameters

ARP/RARP	Any
Request/Reply	Any
Sender IP Filter	Any
Target IP Filter	Any

Request/Reply

Specify the available ARP/RARP opcode (OP) flag for this ACE.

Any: No ARP/RARP 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, enter a specific sender IP address in dotted decimal notation.

Sender IP Mask

When "Network" is selected for the sender IP filter, 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 that appears. **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 that appear.

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.

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").

ARP Sender MAC Match	Any
RARP Target MAC Match	Any
IP/Ethernet Length	Any
IP	Any
Ethernet	Any

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 SMAC address.

1: RARP frames where THA is equal to the SMAC address.

Any: Any value is allowed ("don't-care").

IP/Ethernet Length

Specify if 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 if 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 if 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

The IP parameters can be configured when Frame Type "IPv4" is selected.

IP Protocol Filter

Specify the IP protocol filter for this ACE.

Any: No IP protocol filter is specified ("don't-care").

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.

Other: If you want to filter another specific IP protocol filter with this ACE, choose this value. A field for entering an IP protocol filter displays.

IP Parameters

IP Protocol Filter	Any	▼
IP TTL	Any	▼
IP Fragment	Any	▼
IP Option	Any	▼
SIP Filter	Any	▼
DIP Filter	Any	▼

IP Protocol Value

When "Specific" is selected for the IP protocol value, 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 Parameters

IP Protocol Filter	Other	▼
IP Protocol Value	1	

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.

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 DIP Filter, you can enter a specific DIP address in dotted decimal notation.

DIP Mask

When "Network" is selected for the DIP Filter, you can enter a specific DIP mask in dotted decimal notation.

ICMP Parameters

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 Parameters

ICMP Type Filter	Any	▼
ICMP Code Filter	Any	▼

ICMP Type Value

When "Specific" is selected for the ICMP filter, you can enter a specific ICMP 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

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 Parameters

Source Port Filter	Range
Source Port Range	0 - 65535
Dest. Port Filter	Specific
Dest. Port No.	0
TCP FIN	0
TCP SYN	Any
TCP RST	Any
TCP PSH	Any
TCP ACK	Any
TCP URG	Any

Source Port 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.

Source Port 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.

Destination Port 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 displays.

Dest. Port No.

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.

Destination Port 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

One of several TCP flag names used only when filtering TCP (*urg*, *ack*, *psh*, *rst*, *syn*, and *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

One of several TCP flag names used only when filtering TCP (*urg*, *ack*, *psh*, *rst*, *syn*, and *fin*).

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

One of several TCP flag names used only when filtering TCP (*urg, ack, psh, rst, syn, and fin*).

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

One of several TCP flag names used only when filtering TCP (*urg, ack, psh, rst, syn, and fin*).

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

One of several TCP flag names used only when filtering TCP (*urg, ack, psh, rst, syn, and fin*).

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

One of several TCP flag names used only when filtering TCP (*urg, ack, psh, rst, syn, and fin*).

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").

UDP Parameters

Source Port 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.

UDP Parameters

Source Port Filter	Specific	▼
Source Port No.	0	
Dest. Port Filter	Range	▼
Dest. Port Range	0	65535

Source Port 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.

Source Port 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.

Dest. Port Filter

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.

Destination Port 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.

Ethernet Type Parameters

The Ethernet Type parameters can be configured when Frame Type "Ethernet Type" is selected.

Ethernet Type Parameters

EtherType Filter	Specific ▼
Ethernet Type Value	0xFFFF

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.

Parameter Summary

ACE Configuration

Ingress Port = Port 1, Port 2, Port 3, Port 4
 Policy Filter = Any, Specific
 Policy Value = x (only if Policy Filter = Specific)
 Policy Bitmask = 0x x (only if Policy Filter = Specific)
 Frame Type = Any, Ethernet Type, ARP, IPv4

Action = Permit, Deny
 Rate Limiter = Disabled or 1-16
 EVC Policer = Disabled or Port 1 - Port 5
 Port Redirect = Disabled, Port1, Port 2, Port 3, or Port 4
 Mirror = Enabled or Disabled
 Logging = Enabled or Disabled
 Shutdown = Enabled or Disabled
 Counter = count

Ethernet Type Parameters

Ether Type Filter = Any or Specific
 Ethernet Type Value = ox ffff (only if Ether Type Filter = Specific)

IP Parameters

IP Protocol Filter = Any, ICMP, UDP, TCP, Other

IP TTL = Any, Non-zero, Zero
 IP Fragment = Any, Yes, No
 IP Option = Any, Yes, No
 SIP Filter = Any, Host, Network
 DIP Filter = Any, Host, Network

MAC Parameters

SMAC Filter = Any, Specific
 DMAC Filter = Any, MC, BC, UC

ARP Parameters

ARP/RARP = Any, ARP, RARP, Other
 Request/Reply = Any, Request, Reply
 Sender IP Filter = Any, Host, Network
 Target IP Filter = Any, Host, Network

ARP SMAC Match = Any, 0, 1
 RARP DMAC Match = Any, 0, 1
 IP/Ethernet Length = Any, 0, 1
 IP = Any, 0, 1
 Ethernet = Any, 0, 1

VLAN Parameters

802.1Q Tagged = Any, Disabled, Enabled
 VLAN ID = 2-xxxxx (only if VLAN ID Filter = Specific)
 VLAN ID Filter = Any, Specific
 Tag Priority = 0, 1, 2, 3, 4, 5, 6, 7, Any

ICMP Parameters

ICMP Type Filter = Any or Specific
 ICMP Type Value = 0 - 255
 ICMP Code Filter = Any or Specific
 ICMP Code Value = 0 - 255

UDP Parameters

Source Port Filter = Any, Specific, or Range
 Source Port No. = 0 to 65535
 Source Port Range = 0 to 65535
 Dest. Port Filter = Any, Specific, or Range
 Dest. Port No. = 0 to 65535
 Dest. Port Range = 0 to 65535

TCP Parameters

Source Port Filter = Any, Specific, or Range
 Source Port Range = 0 to 65535
 Dest. Port Filter = Any, Specific, or Range
 Dest. Port No. = 0 to 65535
 TCP FIN = Any or 01
 TCP SYN = Any or 01
 TCP RST = Any or 01
 TCP PSH = Any or 01

TCP ACK = Any or 01

TCP URG = Any or 01

(TCP flag names for use only when filtering TCP; the flag names for the TCP flags are urg, ack, psh, rst, syn, and fin.)

Bandwidth Profile using ACE (Access Control Entry)

Apart from MEF specified layer-2 services, the LIB-44xx can associate the bandwidth profile parameters of < CIR, CBS, EIR, EBS, CM, CF> with any kind of traffic flow. The web interface provides the option of Layer 2 to Layer 4 flows to be associated with a bandwidth profile. The flow can be characterized by different Layer 2-4 options together and rules can be assigned to such a flow. One of the rules can be bandwidth. A sample screen for a TCP flow over a VLAN on Port 2 is shown below:

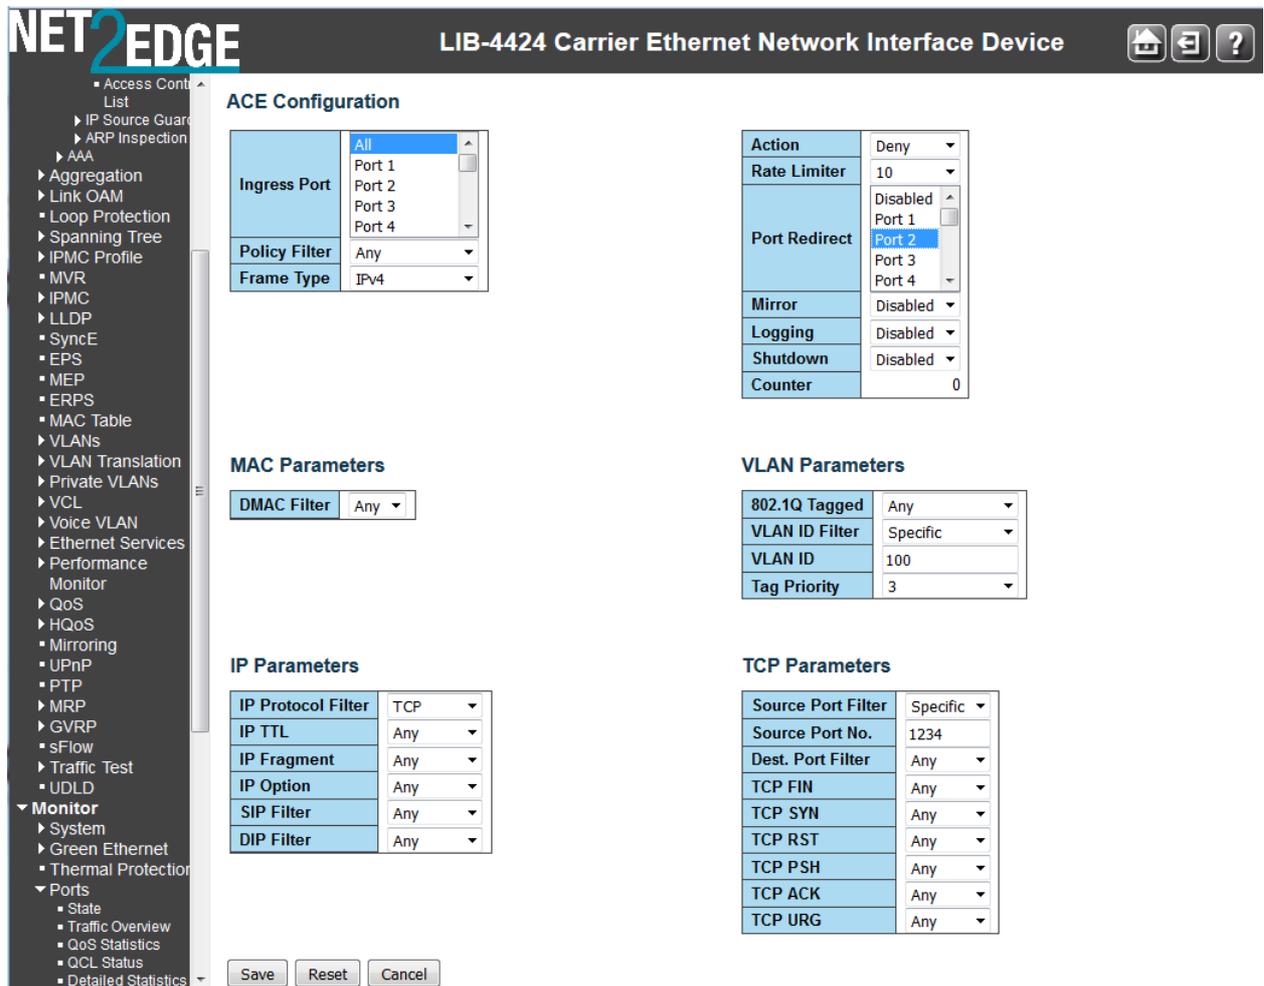
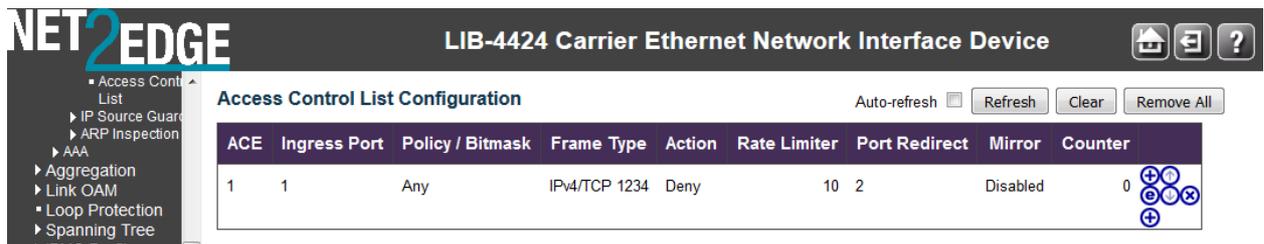


Figure 11: ACE for Flow-based BWP

The resulting ACL Configuration page is shown below:



DHCP Configuration

DHCP (Dynamic Host Configuration Protocol) is used for assigning dynamic IP addresses to devices on a network. DHCP used by networked computers (clients) to obtain IP addresses and other parameters such as the default gateway, subnet mask, and IP addresses of DNS servers from a DHCP server.

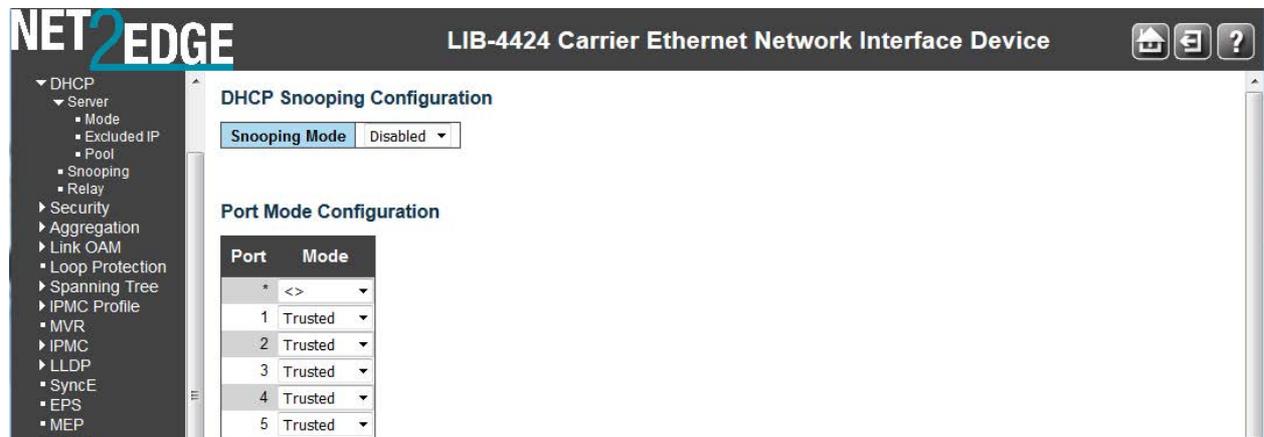
The DHCP server ensures that all IP addresses are unique, for example, no IP address is assigned to a second client while the first client's assignment is valid (its lease has not expired). Therefore, IP address pool management is done by the server and not by a human network administrator.

Dynamic addressing simplifies network administration because the software keeps track of IP addresses rather than requiring an administrator to manage the task. This means that a new computer can be added to a network without the hassle of manually assigning it a unique IP address.

The **Configuration > Security > Network > DHCP** menu path lets you configure LIB-44xx DHCP Snooping and/or DHCP Relay, as discussed below:

DHCP Snooping Configuration

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 DHCP client and server.



The DHCP Snooping parameters are explained below:

Snooping Mode

Indicates 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

The DHCP Snooping feature provides security by inspecting ingress packets for the correct IP and MAC address information. The DHCP Snooping feature defines the LIB-44xx ports as either "trusted" or "untrusted". Enabling DHCP Snooping addresses two network security issues:

All ingress DHCP packets are examined on the untrusted ports and only authorized packets are passed through the switch. Unwanted ingress DHCP packets are discarded.

DHCP ingress packets on an untrusted port are inspected to ensure that the source IP Address and MAC Address combination in each packet is valid when compared to the DHCP Snooping Binding Table. If match is not found, the packet is discarded.

Sets/indicates the DHCP snooping port mode. The valid port modes are:

Trusted: Configures the port as trusted source of the DHCP messages. **Trusted** ports inherently trust all ingress Ethernet traffic. This type of port does no checking or testing on ingress packets. A trusted port connects to a DHCP server either:

Directly to the legitimate trusted DHCP Server

A network device relaying DHCP messages to and from a trusted server

Another trusted source such as a switch with DHCP Snooping enabled.

Untrusted: Configures the port as untrusted source of the DHCP messages. **Untrusted** ports' Ethernet traffic is inherently not trusted. The ingress packets are subsequently tested against specific criteria to determine whether to forward them through the switch, or immediately discard them. Untrusted ports are connected to DHCP clients and to traffic that originates outside of the LAN.

Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

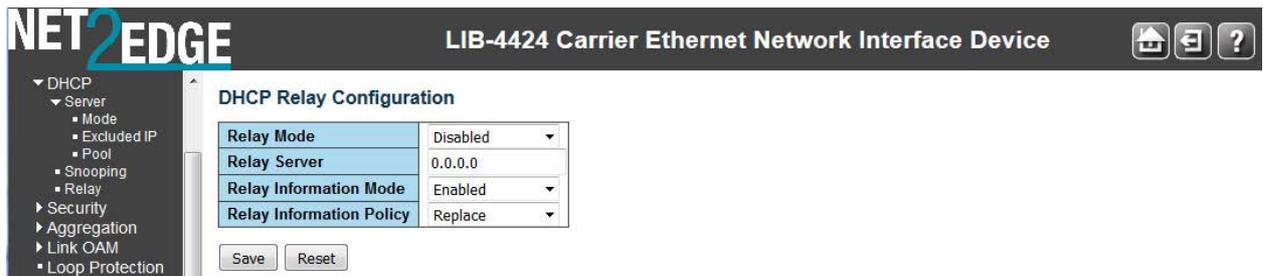
DHCP Relay Configuration

You can configure DHCP Relay from the **Configuration > Security > Network > DHCP > Relay** menu path. 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.

The DHCP option 82 enables a DHCP relay agent to insert specific information into a DHCP request packets when forwarding client DHCP packets to a DHCP server and remove the specific information from a DHCP reply packets when forwarding server DHCP packets to a DHCP client. The DHCP server can use this information to implement IP address or other assignment policies. Specifically the option works by setting two sub-options: Circuit ID (option 1) and Remote ID (option2). The Circuit ID sub-option is supposed to include information specific to which circuit the request came in on. The Remote ID sub-option was designed to carry information relating to the remote host end of the circuit.

The definition of Circuit ID in the switch is 4 bytes long and the format is "*vlan_id*" "*module_id*" "*port_no*". The parameter of "*vlan_id*" is the first two bytes represent the VLAN ID. The "*module_id*" parameter is the third byte for the module ID (in standalone switch it always equal 0). The "*port_no*" parameter is the fourth byte and it means the port number.

The Remote ID is 6 bytes long, and the value is equal to the DHCP relay agents MAC address.



Each of the DHCP Relay Configuration table parameters is explained below:

Relay Mode

Sets / 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

Sets / indicates the DHCP relay server IP address. 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.

Relay Information Mode

Sets / indicates the DHCP relay information mode option operation. The option 82 circuit ID format is "[vlan_id][module_id][port_no]". The first four characters represent the VLAN ID, the fifth and sixth characters are the module ID (on standalone devices it always equal 0), and the last two characters are the port number. For example, "00030108" means the DHCP message was received from VLAN ID 3, switch ID 1, port number 8. And the option 82 remote ID value is equal to the switch MAC address.

The valid 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

Sets / indicates the DHCP relay information option policy. When DHCP relay information mode operation is enabled, if agent receives a DHCP message that already contains relay agent information it will enforce the policy. And it only works under DHCP if relay information operation mode is enabled. Valid 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.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

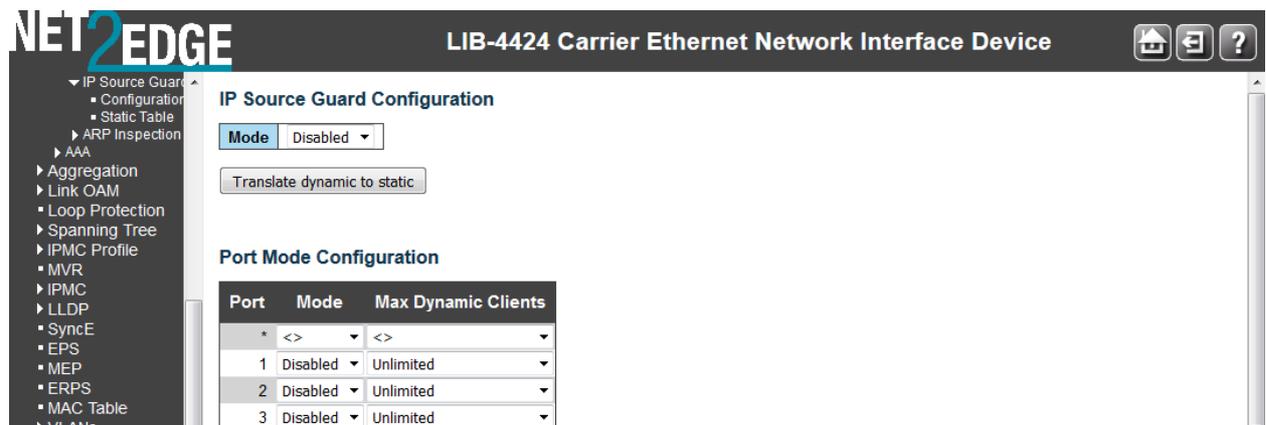
Note: Make sure the DHCP server is connected to a “trusted” LIB-44xx port. See the previous section on “DHCP Snooping Configuration”.

IP Source Guard Configuration

IP Source Guard is a security 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 checks the IP source address and MAC source address in a packet sent from a host attached to an untrusted access interface on the switch against entries stored in the DHCP snooping database. If IP source guard determines that the packet header contains an invalid source IP address or source MAC address, it ensures that the switch does not forward the packet (i.e., the packet is discarded). When you configure IP source guard, you enable on it on one or more VLANs. IP source guard applies its checking rules to packets sent from untrusted access interfaces on those VLANs.

After the DHCP snooping database is populated (via either dynamic DHCP snooping or configuring specific static IP address/MAC address bindings) the IP source guard database is built. It then checks incoming packets from access interfaces on the VLANs on which it is enabled. If the source IP addresses and source MAC addresses match the IP source guard binding entries, the switch forwards the packets to their specified destination addresses. If they do not match, the packets are discarded.

The **Configuration > Security > Network > IP Source Guard** menu path provides Configuration and Static table configuration.



The screenshot displays the configuration page for IP Source Guard. The 'Mode' is currently set to 'Disabled'. A button labeled 'Translate dynamic to static' is visible. The 'Port Mode Configuration' table is as follows:

Port	Mode	Max Dynamic Clients
*	<>	<>
1	Disabled	Unlimited
2	Disabled	Unlimited
3	Disabled	Unlimited

The DHCP Relay Configuration table parameters are explained below:

IP Source Guard > Configuration

Here you can configure IP Source Guard Mode and Port Mode, and translate dynamic to static entries.

Mode

Enable or disable the Global IP Source Guard globally. All configured ACEs will be lost when the mode is enabled globally here.

Port Mode Configuration

Port

The Port column specifies if IP Source Guard is enabled on each port (1-4). The * in the Port column acts as a 'wild card' character which causes the selections in this row to be applied to all other rows in the table for which this selection is valid. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Mode

The table specifies if 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 Clients' is equal to **0**, it means to only allow the IP packets forwarding that are matched in static entries on the specific port.

Buttons

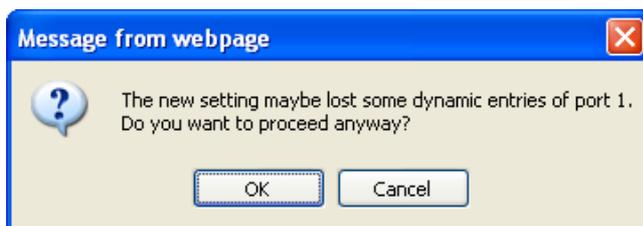
Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Translate dynamic to static: Click to translate all dynamic entries to static entries. The IP Source Guard Configuration Mode must be set to Enabled, and an entry must exist in the Static IP Source Guard Table.

Messages

Message: *The new setting maybe lost some dynamic entries of port 1. Do you want to proceed anyway?*



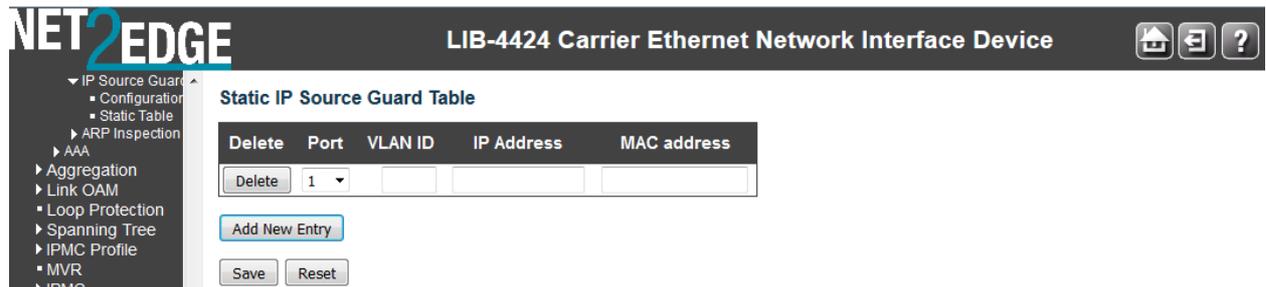
Meaning:

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Change either the "**Mode**" or the "**Max Dynamic Clients**" setting; see above.

IP Source Guard > Static Table

The default Static IP Source Guard Table displays no saved entries. When you click the **Add new entry** button, initial entry fields display.



Delete

Check to delete an existing entry. It will be deleted during the next save.

Port

The logical port for the settings. Select a port from the dropdown (e.g. 1-4 on the LIB-4400). Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

VLAN ID

Enter the VLAN ID for the settings.

IP Address

Enter the allowed Source IP address (e.g., enter *192.168.1.30*).

IP Mask

It can be used for calculating the allowed network with IP address. A valid IP mask is a dotted decimal string (x.y.z.w) where x, y, and z are decimal numbers from 0-255, and when converted to a 32-bit binary string and read left to right, all bits after the first zero must also be zero.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Add New Entry: Click to add a new entry to the 'Static IP Source Guard' table. Specify the Port, VLAN ID, IP address, and IP Mask for the new entry. Click the **Save** button when done.

Example

The Static IP Source Guard Table shown below shows four new saved entries.

Static IP Source Guard Table

Delete	Port	VLAN ID	IP Address	MAC address
<input type="checkbox"/>	1	1	192.168.1.30	00-11-22-33-44-55
<input type="checkbox"/>	1	2	192.168.1.210	00-11-22-33-44-55
<input type="checkbox"/>	2	3	192.168.1.210	00-11-22-33-44-55
<input type="checkbox"/>	3	2	192.168.1.210	00-11-22-33-44-55
<input type="checkbox"/>	3	4	192.168.1.110	00-11-22-33-44-55

The example above has Port 1 with two entries with the same VLAN ID, different IP Addresses, and different IP Masks. Port 3 has two entries for two different VLAN IDs with the same IP Address, and the same IP Mask.

ARP Inspection

ARP Inspection is a security feature. Several types of attacks can be launched against a host or devices connected to Layer 2 networks by "poisoning" the ARP caches. The ARP Inspection feature is used to block such attacks. Only valid ARP requests and responses can go through the switch device.

The **Configuration > Security > Network > ARP Inspection > Configuration** menu path provides global ARP Inspection configuration and Port level ARP Inspection configuration.

NET2EDGE LIB-4424 Carrier Ethernet Network Interface Device

ARP Inspection Configuration

Mode: Disabled

Translate dynamic to static

Port Mode Configuration

Port	Mode	Check VLAN	Log Type
*	<>	<>	<>
1	Disabled	Disabled	None
2	Disabled	Disabled	None
3	Disabled	Disabled	None
4	Disabled	Disabled	None
5	Disabled	Disabled	None

The ARP Inspection parameters are explained below:

ARP Inspection Configuration

Mode

Enable or disable ARP Inspection on a global basis. The default is 'Disabled'.

Port Mode Configuration

Port

Specify if ARP Inspection is enabled on each port. The * in the Port column acts as a 'wild card' character which causes the selections in this row to be applied to all other rows in the table for which this selection is valid. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Mode

Specify ARP Inspection is enabled on which ports. ARP Inspection is enabled on this given port only when both Global Mode and Port Mode on a given port are **Enabled**. The default is '**Disabled**'.

Buttons

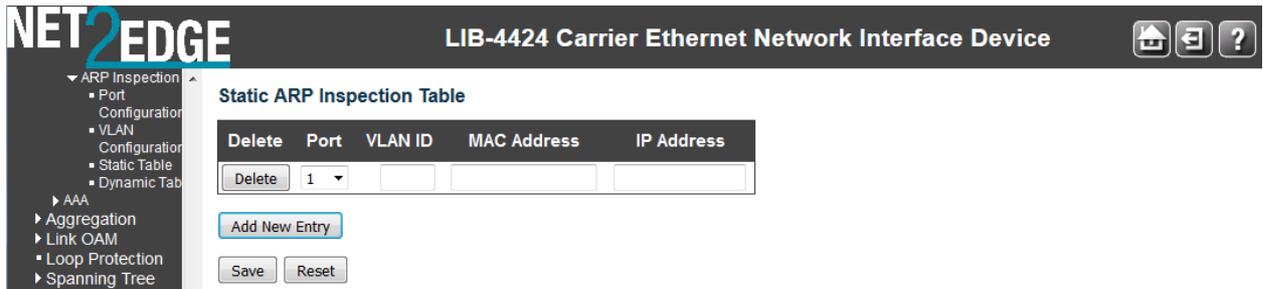
Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Translate dynamic to static: Click to translate all dynamic entries to static entries.

Static ARP Inspection

The **Configuration > Security > Network > ARP Inspection > Static Table** menu path provides the default Static ARP Inspection Table. Click the **Add New Entry** button to display the entry table.



The Static ARP Inspection Table parameters are explained below:

Delete

Check to delete the entry. It will be deleted during the next save.

Port

The logical port for the settings (e.g., ports 1-4 on the LIB-4400). Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

VLAN ID

The VLAN ID for the settings (1-4094).

MAC Address

Allows Source MAC address in ARP request packets (e.g., enter *00-c0-f2-56-08-b0*).

IP Address

Allows Source IP address in ARP request packets.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Add New Entry: Click to add a new entry to the Static ARP Inspection table. Specify the Port, VLAN ID, MAC address, and IP address for the new entry. Click the "Save" button when done adding entries.

Example

The Static ARP Inspection Table shown below shows three saved entries.

The screenshot displays the NET2EDGE web interface for the LIB-4424 Carrier Ethernet Network Interface Device. The left sidebar contains a navigation menu with the following items: ARP Inspection (expanded), Port Configuration, VLAN Configuration, Static Table, Dynamic Table, AAA, Aggregation, Link OAM, Loop Protection, Spanning Tree, IPMC Profile, and MVR. The main content area is titled "Static ARP Inspection Table" and features a table with the following data:

Delete	Port	VLAN ID	MAC Address	IP Address
<input type="checkbox"/>	1	1	00-11-22-33-44-55	192.168.1.30
<input type="checkbox"/>	1	2	00-11-22-33-44-55	192.168.1.210
<input type="checkbox"/>	2	3	00-11-22-33-44-55	192.168.1.210

Below the table are three buttons: "Add New Entry", "Save", and "Reset".

Note that you can create multiple entries and then Save them in one Save operation. You can also cut and paste information between fields. Use the keyboard Tab key to move from one field to the next.

Configuration > Security > AAA

This page lets you configure the AAA (Authentication, Authorization and Accounting) Servers. You can configure the optional RADIUS and/or TACACS+ servers from the **Configuration > Security > AAA** menu path.

RADIUS (Remote Authentication Dial In User Service) networking protocol provides centralized access, authorization and accounting management for computers to connect and use a network service.

TACACS+ (Terminal Access Controller Access Control System Plus) networking protocol 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.

Caution: Before enabling **RADIUS** or **TACACS+**, make sure that the related AAA server is operational and that at least the AAA server's IP address/hostname and encryption/decryption key parameters are set correctly. Make sure using the following safe method:

1. Open a CLI session.
2. Enter the command "**security aaa con**".
3. Try to open a TELNET session.

Now, if the attempt fails (possibly because of an incorrect AAA parameter setting) access to the CLI agent is retained (via the CLI session) and any AAA parameter setting can be corrected in the CLI session.

NET2EDGE LIB-4424 Carrier Ethernet Network Interface Device

AAA

- ▾ RADIUS
- ▾ TACACS+
- Aggregation
- Link OAM
- ▾ Loop Protection
- Spanning Tree
- IPMC Profile
- ▾ MVR
- IPMC
- LLDP
- ▾ SyncE
- ▾ EPS
- ▾ MEP
- ▾ ERPS
- ▾ MAC Table
- VLANs
- VLAN Translation
- Private VLANs
- VCL
- Voice VLAN
- Ethernet Services
- Performance Monitor

RADIUS Server Configuration

Global Configuration

Timeout	5	seconds
Retransmit	3	times
Deadtime	0	minutes
Change Secret Key	No	
NAS-IP-Address		
NAS-IPv6-Address		
NAS-Identifier		

Server Configuration

Delete	Hostname	Auth Port	Acct Port	Timeout	Retransmit	Change Secret Key
Add New Server						
Save Reset						

TACACS+ Server Configuration

Global Configuration

Timeout	5	seconds
Deadtime	0	minutes
Change Secret Key	No	

Server Configuration

Delete	Hostname	Port	Timeout	Change Secret Key
Delete		49		

Buttons: Add New Server, Save, Reset

Common Server Configuration

These settings are common for all of the Authentication Servers (i.e., these parameters apply to both Radius and Tacacs+).

Timeout

The Timeout, which can be set to a number between **3** and **3600** seconds, is the maximum time to wait for a reply from a server. If the server does not reply within this timeframe, it is considered to be 'dead' and continue with the next enabled server (if any).

RADIUS servers use the UDP protocol, which is unreliable by design. In order to cope with lost frames, the timeout interval is divided into 3 subintervals of equal length. If a reply is not received within the subinterval, the request is transmitted again. This algorithm causes the RADIUS server to be queried up to 3 times before it is considered to be dead.

Dead Time

The Dead Time, which can be set to a number from **0** - **3600** seconds, 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 LIB-44xx from continually trying to contact a server that it has already determined as dead. Setting the Dead Time to a value greater than **0** (zero) will enable this feature, but only if more than one server has been configured.

RADIUS Authentication Server Configuration

The table has one row for each RADIUS Authentication Server and a number of columns, which are:

#

The RADIUS Authentication Server number (1-5) for which the configuration below applies.

Enabled

Enable the RADIUS Authentication Server by checking this checkbox.

IP Address/Hostname

The IP address or hostname of the RADIUS Authentication Server. IP address is expressed in dotted decimal notation.

Port

The UDP port to use on the RADIUS Authentication Server. If the port is set to 0 (zero), the default port (**1812**) is used on the RADIUS Authentication Server.

Secret

Enter a secret (password) of up to 29 characters to be shared between the RADIUS Authentication Server and the LIB-44xx.

RADIUS Accounting Server Configuration

The table has one row for each RADIUS Accounting Server and a number of columns, which are:

#

The RADIUS Accounting Server number for which the configuration below applies.

Enabled

Enable the RADIUS Accounting Server by checking this box.

IP Address/Hostname

The IP address or hostname of the RADIUS Accounting Server. IP address is expressed in dotted decimal notation.

Port

The UDP port to use on the RADIUS Accounting Server. If the port is set to **0** (zero), the default port (**1813**) is used on the RADIUS Accounting Server.

Secret

Enter a secret (password) of up to 29 characters to be shared between the RADIUS Accounting Server and the LIB-44xx.

TACACS+ Authentication Server Configuration

The table has one row for each TACACS+ Authentication Server and a number of columns, which are:

#

The TACACS+ Authentication Server number for which the configuration below applies.

Enabled

Enable the TACACS+ Authentication Server by checking this box.

IP Address/Hostname

The IP address or hostname of the TACACS+ Authentication Server. IP address is expressed in dotted decimal notation.

Port

The TCP port to use on the TACACS+ Authentication Server. If the port is set to 0 (zero), the default port (**49**) is used on the TACACS+ Authentication Server.

Secret

Enter a secret (password) of up to 29 characters to be shared between the TACACS+ Authentication Server and the LIB-44xx.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

The **Monitor > Security > AAA** menu path provides **RADIUS Overview** and **RADIUS Details** information.

Aggregation Configuration

Link aggregation (AKA trunking, link bundling, or Ethernet NIC bonding) are various methods of combining (aggregating) multiple network connections in parallel to increase throughput over that which a single connection could sustain, while providing redundancy in case one of the links fails.

Link aggregation addresses two Ethernet connection problems: bandwidth limitations and lack of resilience.

The LIB-44xx supports both Static aggregation and LACP. Note that Aggregation mismatch can occur if the aggregation type on each end of the link does not match. Some switches do not implement the 802.1AX standard but support static link aggregation. Link aggregation between similarly 'statically' configured switches will work, but will fail between a statically configured switch and a device that is configured for LACP.

Link aggregation bundles multiple ports (member ports) together into a single logical link. It is used mainly to increase available bandwidth without introducing loops into the network, and to improve resilience against faults. A link aggregation group (LAG) can be established with individual links being dynamically added or removed. This enables bandwidth to be incrementally scaled based on changing requirements. A LAG can be quickly reconfigured if faults are identified.

Frames destined for a LAG are sent on only one of the LAGs member ports. The member port on which a frame is forwarded is determined by a 4-bit aggregation code (AC) that is calculated for the frame.

The aggregation code ensures that frames belonging to the same frame flow (e.g., a TCP connection) are always forwarded on the same LAG member port. For that reason, reordering of frames within a flow is not possible. The AC is based on the following information:

- SMAC (Source MAC address)
- DMAC (Destination MAC address)
- Source and Destination IPv4 address
- Source and Destination TCP/UDP ports for IPv4 packets
- Source and Destination TCP/UDP ports for IPv6 packets
- IPv6 Flow Label

For best traffic distribution among LAG member ports, enable all six contributions to the AC.

Each LAG can consist of up to 16 member ports. Any quantity of LAGs may be configured for the LIB-44xx (only limited by the number of device ports). To configure a proper traffic distribution, the ports within a LAG must use the same link speed.

A port cannot be a member of multiple LAGs.

The Aggregation Configuration parameters are explained below:

Static Aggregation

This page is used to configure the Aggregation hash mode and the aggregation group from the **Configuration > Aggregation > Static** menu path.

Aggregation involves using multiple ports in parallel to increase the link speed beyond the limits of a port, and to increase the redundancy for higher availability (aka 'Port Aggregation' or 'Link Aggregation').

Guidelines for static aggregation:

A static aggregation can contain up to five ports.

Before creating a port aggregation, verify that the settings are the same for all ports in the trunk including speed, duplex mode, flow control, VLAN membership, etc. If these settings are not the same, the switch will not let you create the aggregation.

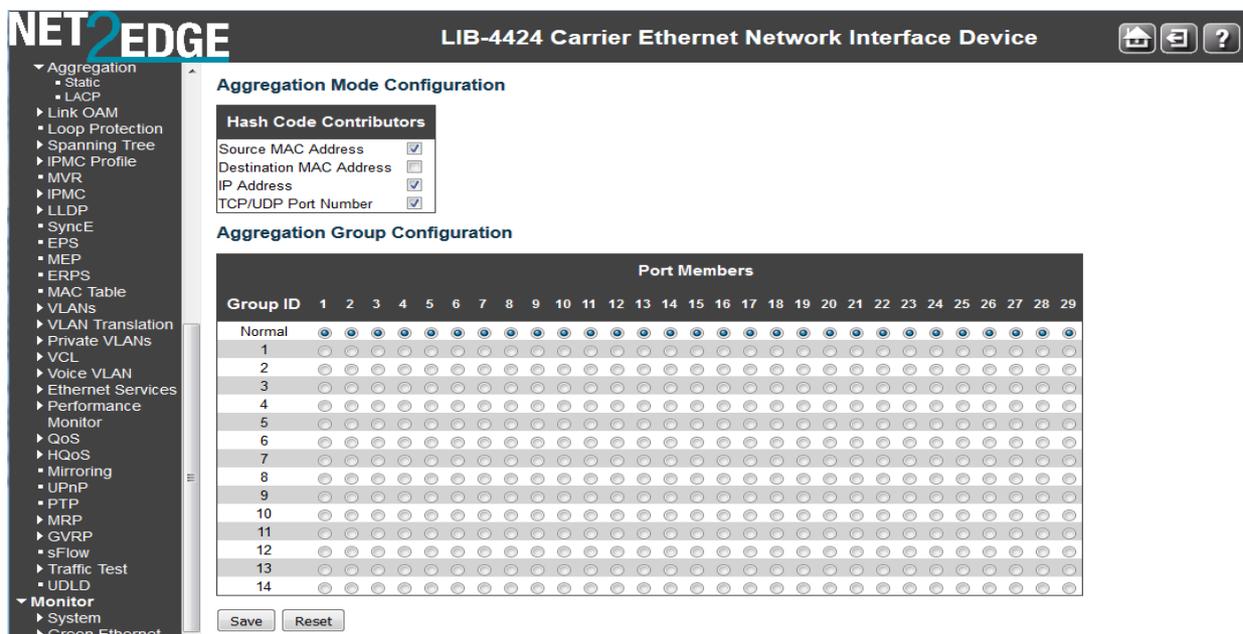
After you create a port aggregation, a change to the speed, duplex mode, flow control, or back pressure of any port in the trunk automatically implements the same change on all the other member ports.

A port can belong to only one static aggregation at a time.

The ports of a static aggregation can be configured to be members of more than one VLAN.

The ports of a static aggregation can be either tagged or untagged members of the same VLAN.

The default **Configuration > Aggregation > Static** page is shown below:



The static aggregation parameters are explained below:

Aggregation Mode Configuration - Hash Code Contributors

Aggregation is possible based on hash values computed from fields in the frame.

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 as a hash code contributor.

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 as a hash code contributor.

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 as a hash code contributor.

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 as a hash code contributor.

Aggregation Group Configuration

Group ID

Indicates the 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 LIB-44xx 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 of the same speed in each group. Each local aggregation group must contain 2 - 16 members.

Static aggregation cannot be enabled on ports whose 802.1X Admin State is not 'Authorized'. To configure the Spanning Tree function, see **Configuration > Spanning Tree > CIST Ports > CIST Normal Port Configuration** in the "STP Enabled" column.

Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Messages:

Aggregation Error - Port joining aggregation must be in the same speed and in full duplex

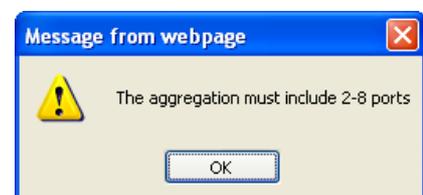
Aggregation Error - Static aggregation cannot be enabled on ports whose 802.1X Admin State is not Authorized

Group 1 member counts error!! - Local aggregation must include 2-4 ports.

LACP Error - LACP and Static aggregation cannot both be enabled on the same ports

The aggregation must include 2-4 ports

The aggregation must include 2-8 ports.



Set the aggregation equal to or greater than 2 ports.

LACP (Link Aggregation Control Protocol)

This page is used to configure the [Aggregation](#) hash mode and the aggregation group from **Configuration > Aggregation > LACP**. The Link Aggregation Control Protocol (LACP) is an IEEE 802.3ad standard that allows bundling several physical ports together to form a single logical port.

The LIB-44xx supports Link aggregation per IEEE 802.1AX-2008. The Link aggregation supports several physical links bundled into a single logical link for resiliency and load sharing. The LIB-44xx uses LACP PDUs to negotiate with peer devices and to exchange information about the links to be bundled automatically when enabled on the physical port.

The resolved aggregation status and peer information status are available. The load sharing mechanism uses all the physical links to transfer the traffic, but for a flow only one link can be used to make sure the packets are sent/received in order. It uses a hash function to determine which port should carry a traffic flow. The device lets you choose the fields that are needed for generating the hash code needed for routing a flow through a single physical port belonging to the aggregate group.

LACP takes care of link failures where if one link fails the flows belonging to that link are transferred to another link based on the hash mechanism which needs to choose from the available links. The static aggregation option is also supported so the LIB-44xx will work with devices which don't support LACP.

LACP works by sending frames (LACPDUs) down all links that have the protocol enabled. If it finds a device on the other end of the link that also has LACP enabled, it will also independently send frames along the same links enabling the two units to detect multiple links between themselves and then combine them into a single logical link. LACP can be configured in one of two modes: active or passive. In active mode it will always send frames along the configured links. In passive mode however, it acts as "speak when spoken to", and therefore can be used as a way of controlling accidental loops (as long as the other device is in active mode).

LACP has advantages over static configuration. For instance, with failover when a link fails and there is another device such as a Media Converter between the devices, which means that the peer will not see the link down. With static link aggregation the peer would continue sending traffic down the link causing it to be lost. The device can confirm that the configuration at the other end can handle link aggregation. With Static link aggregation, a cabling or configuration mistake could go undetected and cause undesirable network behaviour.

Guidelines for creating LACP aggregations:

The LACP module can have a maximum of 4 groups (number of ports / 2), and up to 5 ports can be in a LAG (Link Aggregation Group) at any time.

LACP must be activated on both the LIB-44xx and its partner device.

The other device must be 802.3ad-compliant.

The ports of an aggregate trunk must be the same medium type (all TP ports or all fibre ports).

Aggregation ports can be consecutive or non-consecutive.

A port can belong to only one LACP aggregator at a time.

A port cannot be a member of an LACP aggregator and a static aggregation concurrently.

The ports of an aggregate trunk must be untagged members of the same VLAN.

LACP trunking is not supported in half-duplex mode. Twisted-pair ports must be set to Auto-Negotiation or 1000 Mbps / full-duplex mode.

1000Base-X fibre optic ports must be set to full-duplex mode.

Only ports that are members of an aggregator will transmit LACPDU packets.

A member port of an aggregator functions as part of an aggregate trunk only if it receives LACPDU packets from the remote device. If it does not receive LACPDU packets, it functions as a normal Ethernet port (forwarding network traffic while continuing to transmit LACPDU packets).

The port with the highest priority in an aggregate trunk carries broadcast packets and packets with an unknown destination.

Before creating an aggregate trunk between a device and another vendor’s device, refer to the vendor’s documentation to determine the maximum number of active ports the device can support in a trunk. If less than eight, the maximum number for the LIB-44xx, assign the other vendor’s device a higher system LACP priority than your LIB-44xx. This helps avoid a conflict between the devices if some ports are put in standby mode when the devices create the trunk.

The **Configuration > Aggregation > LACP** menu path is used to configure the Aggregation hash mode and the aggregation group from the LACP Port Configuration table. **Note:** LACP and Static aggregation cannot both be enabled on the same ports at the same time.

Port	LACP Enabled	Key	Role	Timeout	Prio
*	<input checked="" type="checkbox"/>	<>	<>	<>	32768
1	<input checked="" type="checkbox"/>	Auto	Active	Fast	32768
2	<input checked="" type="checkbox"/>	Auto	Active	Fast	32768
3	<input checked="" type="checkbox"/>	Auto	Active	Fast	32768
4	<input checked="" type="checkbox"/>	Auto	Active	Fast	32768
5	<input checked="" type="checkbox"/>	Auto	Active	Fast	32768
6	<input checked="" type="checkbox"/>	Auto	Active	Fast	32768

This page lets you view and modify the following LACP port parameters.

Port

The LIB-44xx port number (e.g., 1-4 for the LIB-4400). The * in the Port column acts as a ‘wild card’ character which causes the selections in this row to be applied to all other rows in the table for which this selection is valid. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

LACP Enabled

Controls whether LACP is enabled on this LIB-44xx port. LACP will form an aggregation when 2 or more ports are connected to the same partner. The LIB-44xx can support up to four LAGs (assuming two ports in each LAG). LACP cannot be enabled on ports whose 802.1X Admin State is not Authorized.

To configure the Spanning Tree function, use the **Configuration > Spanning Tree > CIST Ports > CIST Normal Port Configuration** menu path in the “STP Enabled” column.

Key

The **Key** value incurred by the port, range 1-65535 . The **Auto** setting will set the key as appropriate by the physical link speed, 10Mb = 1, 100Mb = 2, 1Gb = 3. Using the **Specific** setting, a user-defined value can be entered. Ports with the same **Key** value can participate in the same aggregation group, while ports with different Keys cannot.

Role

The **Role** shows the LACP activity status. The **Active** will transmit LACP packets each second, while **Passive** will wait for a LACP packet from a partner (i.e., ‘speak if spoken to’).

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. The default is **Fast**.

Prio

The **Prio** controls the priority of the port. 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. A lower number means greater priority. The default is **32768**. The valid range is **1 – 65,535**.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Messages:

LACP Error - LACP cannot be enabled on ports whose 802.1X Admin State is not Authorized

Static aggregation and LACP cannot both be enabled on the same ports

The aggregation must include 2-8 ports (2-4 ports on the LIB-4400). Click the browser back button and change the configuration to a valid one.

Link OAM (LOAM) Configuration

The Ethernet Operations, Administration and Maintenance (OAM) protocol is used for monitoring and troubleshooting Ethernet networks. The LIB-44xx supports the Link OAM standard IEEE 802.3–2005 Clause 57 which defines mechanisms for monitoring and troubleshooting Ethernet access links. Specifically it defines tools for discovery, remote failure indication, remote and local loopback, and status and performance monitoring. The LIB-44xx supports Link OAM instance at every port. You can enable Link OAM on any port(s). Consider the following when configuring LOAM on the LIB-44xx:

By default LOAM is disabled and set to Passive mode.

The LOAM unidirectional link feature is not supported.

The LOAM discovery process is interoperable with other vendor devices and should be able to perform LOAM loopback and event notifications with other vendor devices.

The LOAM loopback is supported on all ports and can perform line rate loopback of all the data plane traffic only.

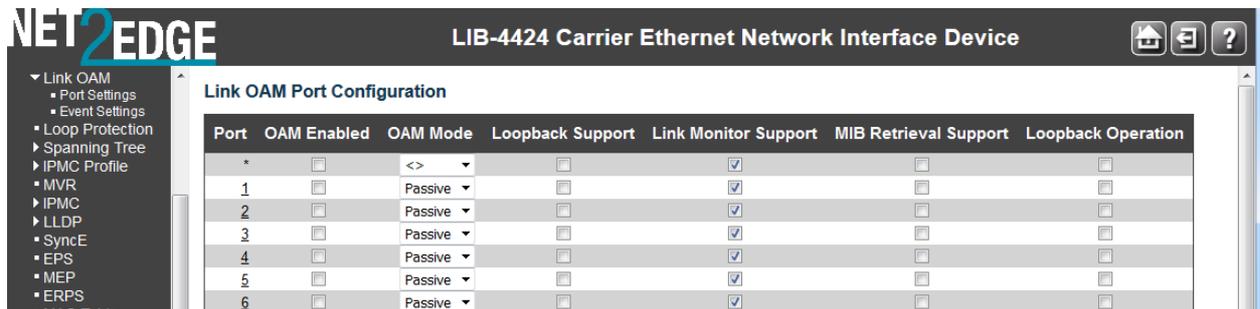
The LOAM Dying gasp is supported; if all ports have LOAM operational, then the priority of sending a dying gasp will be over the uplink ports.

Link events include Errored Symbol Period, Errored Frame Event, and Errored Frame Period Event. The LOAM counter statistics are available via all management interfaces, and an option to reset the LOAM counters is also available.

SNMP traps are generated for dying gasp events.

LIB-44xx Link OAM configuration involves 'Port Settings' and 'Event Settings' as explained below:
Port Settings

This page lets you view and edit the current Link OAM port configurations from the **Configuration > Link OAM > Port Settings** menu path.



Port

The LIB-44xx port number. Click on a specific port number in the table to display that port's 'Detailed Link OAM Status'. The * in the Port column acts as a 'wild card' character which causes the selections in this row to be applied to all other rows in the table for which this selection is valid. See the **Monitor > Link OAM > Port Statistics** section for more information. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

OAM Enabled

Controls whether Link OAM is enabled on this LIB-44xx port. Enabling Link OAM provides the network operator 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: DTEs configured in Active mode initiate the exchange of Information OAMPDUs as defined by the Discovery process. Once the Discovery process completes, Active DTEs are permitted to send any OAMPDU while connected to a remote OAM peer entity in Active mode. Active DTEs 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: DTEs configured in Passive mode do not initiate the Discovery process. Passive DTEs react to the initiation of the Discovery process by the remote DTE. This eliminates the possibility of passive to passive links. Passive DTE' will not send Variable Request or Loopback Control OAMPDUs.

Loopback Support

Controls whether the loopback support is enabled for the LIB-44xx port. Link OAM remote loopback can be used for fault localization and link performance testing. Enabling the loopback support allows the DTE to execute the remote loopback command that helps in fault detection.

Link Monitor Support

Controls whether the Link Monitor support is enabled for the LIB-44xx 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 LIB-44xx port. On enabling the MIB retrieval support, the DTE supports polling of various Link OAM based MIB variables' contents. Note that if MIB Retrieval Support is enabled, 'Loopback Operation' must be disabled. If both are enabled, the message "OAM Error - Error while configuring the OAM loopback" displays. To recover, click the browser's Back button and disable either MIB Retrieval Support or Loopback Operation.

Loopback Operation

If "Loopback Support" is enabled (see above), checking this checkbox will start a loopback operation for the port.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Messages:

Message: OAM Error - Error while configuring the OAM loopback.

Meaning: A Link OAM port configuration is in error (illegal configuration).

Recovery:

1. Click the browser back button.
2. Change the Link OAM Port Configuration (e.g., uncheck the "Loopback Operation" checkbox) and click the **Save** button.

Event Settings

This page lets you view and edit the current Link OAM port configurations from the **Configuration > Link OAM > Event Settings** menu path.

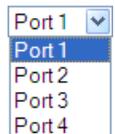
The screenshot shows the 'Link Event Configuration for Port 1' page. The table below represents the data shown in the interface:

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

The Link Event Configuration parameters are explained below:

Port x

The LIB-44xx port number. The port select box determines which port is affected by configuring this page and clicking its buttons. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.



Event Name

Name of the Link Event which is being configured.

Error Window

Represents the window period in the order of **1** second 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** second). 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 from **1** - **60** and its default value is '**1**'. Error Threshold must be between **0** - **0xffffffff** and its default value is '**0**'.

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 from **1** - **60** and its default value is '**1**'. Error Threshold must be between **0** - **0xffffffff** and its default value is '**0**'.

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 from **10** - **900** and its default value is **60**. The Error Threshold must be between **0** - **0xffff** and its default value is **1**.

Buttons

Port 1 . The port select box determines which port is affected by clicking the buttons.

Auto-refresh: Check this checkbox to enable an automatic refresh of the page at 3 second intervals.

Refresh: Click to refresh the page immediately.

Clear: Clears the counters for the selected port.

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Messages: *OAM Error - Error While Configuring Link Events*

Detailed Link OAM Status

After a 'Save', at the **Link OAM Port Configuration** page, click on a Port number in the "Port" column to display that particular port's "Detailed Link OAM Status" information.

The screenshot shows the 'Detailed Link OAM Status for Port 1' page. At the top, there is a navigation menu on the left and a header with 'LIB-4424 Carrier Ethernet Network Interface Device'. The main content area is divided into two columns: 'Local' and 'Peer'. Each column has a table of status indicators.

Local		Peer	
Mode	Passive	Mode	-----
Unidirectional Operation Support	Disabled	Unidirectional Operation Support	-----
Remote Loopback Support	Disabled	Remote Loopback Support	-----
Link Monitoring Support	Enabled	Link Monitoring Support	-----
MIB Retrieval Support	Disabled	MIB Retrieval Support	-----
MTU Size	1500	MTU Size	-----
Multiplexer State	Forwarding	Multiplexer State	-----
Parser State	Forwarding	Parser State	-----
Organizational Unique Identification	00-01-c1	Organizational Unique Identification	-----
PDU Revision	0	PDU Revision	-----

This page provides Link OAM configuration operational status. The displayed fields show the active configuration status for the selected port.

PDU Permission

Displays the port's current level of PDU permissions (e.g., **Info exchange**, **Receive only**).

Discovery State

Displays the port's current state of discovery (e.g., **Active state**, **Fault state**).

Peer MAC Address

Displays the peer's MAC address or "-----" if no peer is available.

Local and Peer

Mode

The Mode in which the Link OAM is operating; **Active** or **Passive**.

Unidirectional Operation Support

This feature is not user configurable. The status of this configuration is retrieved from the PHY (e.g., **Disabled** or **Disabled**).

Remote Loopback Support

If status is enabled, DTE is capable of OAM remote loopback mode (e.g., **Disabled** or **Disabled**).

Link Monitoring Support

If status is enabled, DTE supports interpreting Link Events (e.g., **Disabled** or **Disabled**).

MIB Retrieval Support

If status is enabled, DTE supports sending Variable Response OAMPDUs (e.g., **Disabled** or **Disabled**).

MTU Size

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 (e.g., **1500**).

Multiplexer State

When **Forwarding** displays, the device is forwarding non-OAMPDUs to the lower sublayer. When **Discarding** displays, the device discards all the non-OAMPDUs.

Parser State

When **Forwarding** displays, the device is forwarding non-OAMPDUs to higher sublayer.
When **Loopback**, displays, the device is looping back non-OAMPDUs to the lower sublayer.
When **Discarding** displays, the device is discarding non-OAMPDUs.

Organizational Unique Identification

Displays the 24-bit Organizationally Unique Identifier (OUI) of the vendor, if available (e.g., **00-01-c1**).

PDU Revision

It indicates the current revision of the Information TLV. The value of this field starts 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 does not need to be parsed as nothing in it has changed).

Buttons

. The **Port select box** determines which port is affected by edits on this page.

Refresh: Click to refresh the page immediately.

Auto-refresh: Check this checkbox to enable an automatic refresh of the page at 3 second intervals. Note that the **Monitor > Link OAM** menu path provides LOAM **Statistics**, **Port Status**, and **Event Status** information.

Loop Protection Configuration

The **Configuration > Loop Protection** menu path lets you view and/or change the current global and port-level Loop Protection configuration.

Note: If you will be using the LIB-44xx Loop Protection function, enable Loop Protection here, both globally and at the port level, as one of the first overall configuration steps.

Loop protection increases the efficiency of STP, RSTP, and MSTP by preventing ports from going into a forwarding state that would result in a loop opening up in the network. In spanning tree topologies, a loop-free network is supported by the exchange of a BPDU. Peer STP applications running on the switch interfaces use BPDUs to communicate. The exchange of BPDUs ultimately determines which interfaces block traffic (preventing loops) and which interfaces become root ports and forward traffic. However, a blocking interface can transition to the forwarding state erroneously if the interface stops receiving BPDUs from its designated port on the segment. This transition error can occur with a hardware error on the switch or a software configuration error between the switch and its neighbour.

With loop protection enabled, the spanning tree topology detects root ports and blocked ports, and ensures that both keep receiving BPDUs. If a loop protection enabled interface quits receiving BPDUs from its designated port, it reacts as it would react to a problem with the physical connection on this interface. Rather than transition the interface to a forwarding state, it instead transitions it to a 'loop inconsistent' state. The interface recovers, and then it transitions back to the spanning tree blocking state when it receives a BPDU.

Loop protection is most effective when enabled in the entire switched network. You should generally enable loop protection on all switch interfaces that could become a root or designated port. If you will be using the Loop Protection function, enable Loop Protection here, both globally and at the port level, as one of the first overall configuration steps.

The default Loop Protection screen is shown below:

The screenshot shows the NET2EDGE web interface for a LIB-4424 Carrier Ethernet Network Interface Device. The main content area is titled 'Loop Protection Configuration' and is divided into two sections: 'General Settings' and 'Port Configuration'.

General Settings:

- Global Configuration:**
 - Enable Loop Protection:** Disable (dropdown)
 - Transmission Time:** 5 seconds
 - Shutdown Time:** 180 seconds

Port Configuration:

Port	Enable	Action	Tx Mode
*	<input checked="" type="checkbox"/>	<>	<>
1	<input checked="" type="checkbox"/>	Shutdown Port	Enable
2	<input checked="" type="checkbox"/>	Shutdown Port	Enable
3	<input checked="" type="checkbox"/>	Shutdown Port	Enable
4	<input checked="" type="checkbox"/>	Shutdown Port	Enable
5	<input checked="" type="checkbox"/>	Shutdown Port	Enable

The Loop Protection parameters are explained below:

General Settings

Enable Loop Protection

Controls whether loop protections is enabled (as a whole) or disabled.

Transmission Time

The interval between each loop protection PDU sent on each port. Valid values are **1** - **10** seconds.

Shutdown Time

The period (in seconds) for which a port will be kept disabled if 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 (**0**) will keep a port disabled. The default is 180 seconds (3 minutes). Note that a loop port can be link uped by disabling Loop Protection with "Shutdown Time" set to 0 (zero).

Port Configuration

Port

The switch port number of the port. Note that loop protection is not supported on the MGMT port. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Enable

Controls whether loop protection is enabled on this LIB-44xx port.

Action

Configures the action to be performed when a loop is detected on a port. The valid Actions are:

Sets the action to be performed when a loop is detected on a port. Valid loop protect Action values are:

Shutdown Port: Shutdown the port.

Shutdown and Log : Shutdown the port and Log the event.

Log Only: Only Log the event.

Trap Only: Only send a trap.

Shutdown and Trap: Shutdown the port and Send trap.

Log and Trap: Send Trap and Log the event.

All: Shutdown the port, send trap, and Log the event.

Shutdown Port
Shutdown and Log
Log Only
Trap Only
Shutdown and Trap
Log and Trap
All

Tx Mode

Controls whether the port is actively generating loop protection PDUs, or whether it is just passively looking for looped PDUs.

Enable: this port is actively generating loop protection PDUs.

Disable: this port is passively looking for looped PDUs.

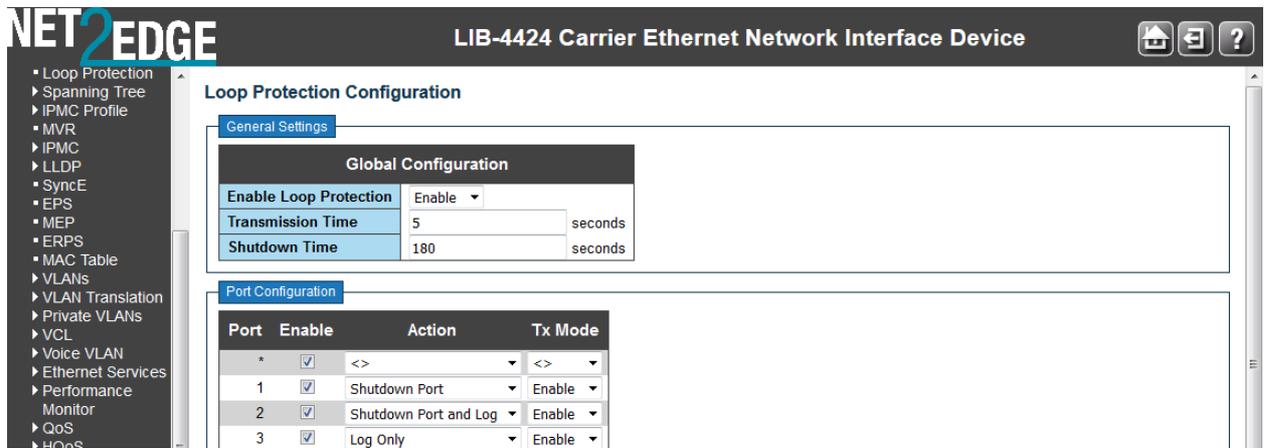
Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Example

A valid, saved Loop Protection configuration example is shown below:



Spanning Tree

The LIB-44xx Spanning Tree menu provides the STP (Spanning Tree Protocol) configuration sub-menus from the **Configuration > Spanning Tree** menu path.

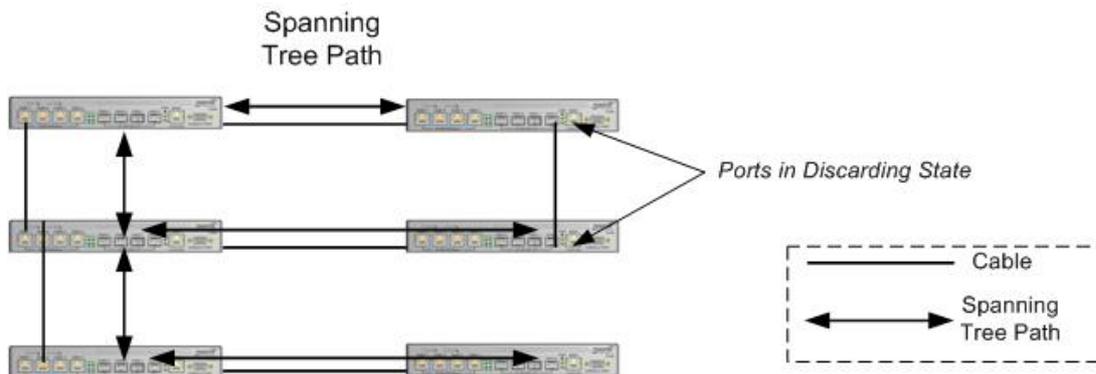


Figure 2. Spanning Tree Example

The Spanning Tree Protocol is an OSI layer-2 protocol which ensures a loop-free topology for any bridged LAN.

STP/RSTP/MSTP

The LIB-44xx supports the spanning tree protocols of STP/RSTP and MSTP on all interfaces. The Spanning Tree protocols help in creating a loop free bridged network. The implementation conforms to the IEEE specs 802.1D for STP, 802.1w for RSTP and 802.1s for MSTP.

The LIB-44xx can act in the role of a root bridge or as a designated bridge by the process of election. The priorities for the bridge instance that is used in BPDU frames can be configured. For MSTP, each MSTI (Multiple Spanning Tree Instance) priority can be configured for the Common and Internal Spanning Tree (CIST) instance.

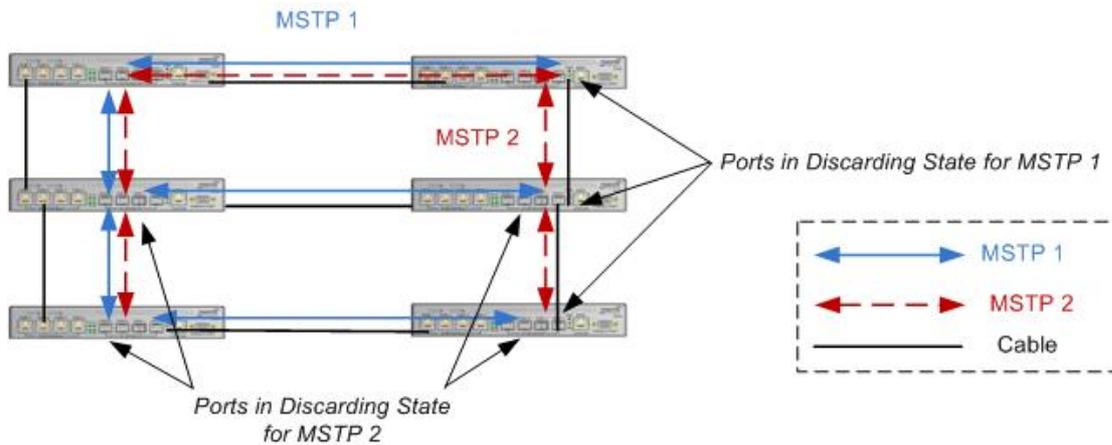


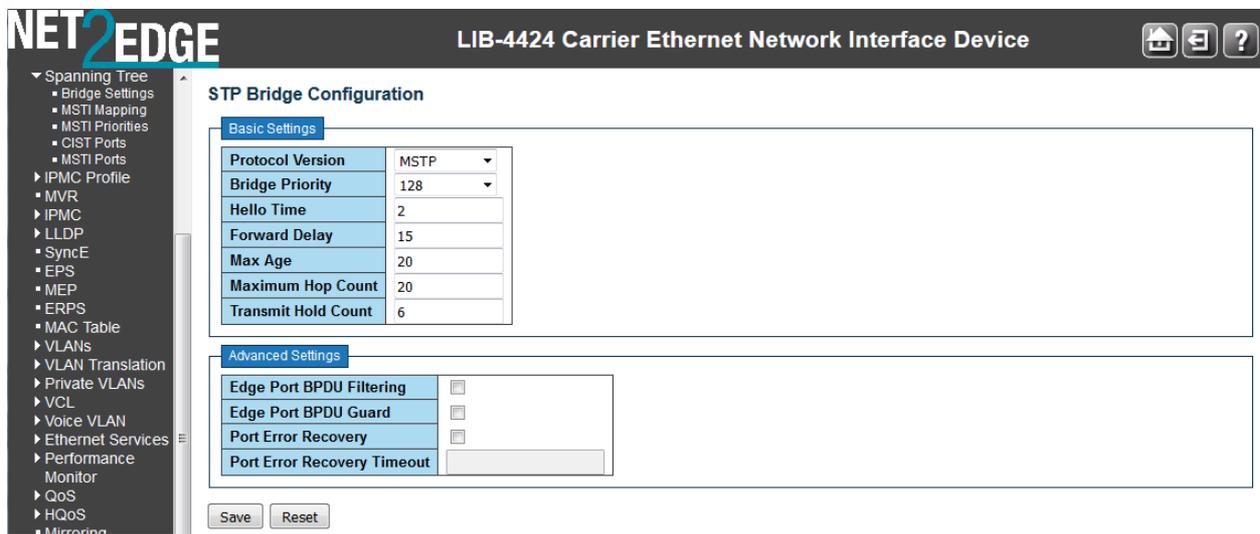
Figure 3. Multiple Spanning Tree Example

The MSTP protocol version works over VLAN instances and multiple VLANs can be added to an MSTI; however, at any time a VLAN can only be part of one MSTI. Configuration for each MSTI and the VLANs that belong to that instance is supported. The LIB-44xx also supports configuration of enabling/disabling BPDU guard, path cost for that port, restricting topology change notification, etc. Note that MSTP is disabled on the LIB-44xx MGMT port.

The Spanning Tree sub-menus (Bridge Settings, MSTI Mapping, MSTI Priorities, CIST Ports, and MSTI Ports) are described in the following sections:

Bridge Settings

LIB-44xx STP Bridge configuration is done from the **Configuration > Spanning Tree > Bridge Settings** menu path.



This page lets you configure STP system settings, which are used by all LIB-44xx STP Bridge instances.

Basic Settings

Protocol Version

The STP protocol version setting. Valid values are **STP**, **RSTP** and **MSTP**.

Bridge Priority

Sets the bridge priority (the priority setting among other switches in the Spanning Tree). Lower numeric values have better priority. The bridge priority plus the MSTI instance number, concatenated with the 6-byte MAC address of the LIB-44xx forms a *Bridge Identifier*.

For **MSTP** operation, this is the priority of the CIST (Common and Internal Spanning Tree).

For **STP** or **RSTP** operation, this is the priority of the STP/RSTP bridge.

Select a Bridge Priority of 0, 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, or 61440. The default is **32768**.

Forward Delay

The delay used by STP Bridges to transit Root and Designated Ports to Forwarding (used in STP compatible mode). Valid values are in the range **4** to **30** seconds. The default is **15** seconds.

Max Age

The maximum age of the information transmitted by the Bridge when it is the Root Bridge. Valid values are

6 - **40** seconds, and **Max Age** must be $\leq (\text{FwdDelay}-1)*2$. The default is **20** seconds.

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 **6** to **40** hops. The default is **20** 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. The default is **6** BPDU's / second.

Advanced Settings

Edge Port BPDU Filtering

Control whether a port which is explicitly configured as **Edge** will transmit and receive BPDUs.

Edge Port BPDU Guard

Check the checkbox to force a port which is 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

Check the checkbox to force a port in the *error-disabled* state to be automatically enabled after a certain time. If recovery is not enabled, ports must be disabled and then 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 **30** seconds to **86400** seconds (24 hours). The 'Port Error Recovery' checkbox (above) must be checked to be able to make an entry in this field.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

MSTI Mapping

MSTI (Multiple Spanning Tree Instance) configuration is done from the **Configuration > Spanning Tree > MSTI Mapping** menu path.

MSTP enables the grouping and mapping of VLANs to different spanning tree instances. An MSTI (MST Instance) is a particular set of VLANs that use the same spanning tree.

NET2EDGE LIB-4424 Carrier Ethernet Network Interface Device

MSTI Configuration

Add VLANs separated by spaces or comma.
Unmapped VLANs are mapped to the CIST. (The default bridge instance).

Configuration Identification

Configuration Name	00-01-c1-00-fc-b0
Configuration Revision	0

MSTI Mapping

MSTI	VLANs Mapped
MSTI1	
MSTI2	
MSTI3	
MSTI4	
MSTI5	
MSTI6	
MSTI7	

Save Reset

This page lets you view and/or edit the current STP MSTI bridge instance priority configurations.
Configuration Identification

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 MSTIs (Intra-region). Enter a name of up to 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.

VLANs Mapped

The list of VLANs mapped to the MSTI. The VLANs 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 have any VLANs mapped to it). Enter a single VLAN ID or a range of VLAN IDs. Note that VLAN **0** is invalid. Any VLAN can only be mapped to one MSTI.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

A sample MSTI Mapping > MSTI Configuration is shown below:

NET2EDGE LIB-4424 Carrier Ethernet Network Interface Device

MSTI Configuration

Add VLANs separated by spaces or comma.

Unmapped VLANs are mapped to the CIST. (The default bridge instance).

Configuration Identification

Configuration Name	00-01-c1-00-fc-b0
Configuration Revision	0

MSTI Mapping

MSTI	VLANs Mapped
MSTI1	1
MSTI2	2, 5
MSTI3	3-4
MSTI4	6-12
MSTI5	
MSTI6	
MSTI7	

Save Reset

MSTI Priorities

MSTI Priority Configuration is done from the **Configuration > Spanning Tree > MSTI Mapping** menu path.

MSTI	Priority
*	<>
CIST	128
MSTI1	128
MSTI2	128
MSTI3	128
MSTI4	128
MSTI5	128
MSTI6	128
MSTI7	128

This page lets you view and/or edit the current STP MSTI bridge instance priority configurations.

MSTI

The bridge instance. The CIST is the *default* instance, which is always active. The * in the MSTI column acts as a 'wild card' character which causes the selections in this row to be applied to all other rows in the table for which this selection is valid.

Priority

Sets the bridge priority. Lower numeric values have better priority. The bridge priority plus the MSTI instance number, concatenated with the 6-byte LIB-44xx MAC address forms a *Bridge Identifier*. Select 0, 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, or 61440. The default is **32768**.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

CIST Ports

CIST (Common and Internal Spanning Tree) Port configuration is done from the **Configuration > Spanning Tree > CST Ports** menu path.

The CIST is the default spanning tree instance of MSTP (i.e., all VLANs that are not members of particular MSTIs are members of the CIST). Also, an individual MST region can be regarded a single virtual bridge by other MST regions. The spanning tree that runs between regions is the CIST.

NET2EDGE LIB-4424 Carrier Ethernet Network Interface Device

STP CIST Port Configuration

CIST Aggregated Port Configuration

Port	STP Enabled	Path Cost	Priority	Admin Edge	Auto Edge	Restricted Role	TCN	BPDU Guard	Point-to-point
-	<input type="checkbox"/>	Auto	128	Non-Edge	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Forced True

CIST Normal Port Configuration

Port	STP Enabled	Path Cost	Priority	Admin Edge	Auto Edge	Restricted Role	TCN	BPDU Guard	Point-to-point
*	<input type="checkbox"/>	<>	<>	<>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<>
1	<input type="checkbox"/>	Auto	128	Non-Edge	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Auto
2	<input type="checkbox"/>	Auto	128	Non-Edge	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Auto
3	<input type="checkbox"/>	Auto	128	Non-Edge	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Auto
4	<input type="checkbox"/>	Auto	128	Non-Edge	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Auto
5	<input type="checkbox"/>	Auto	128	Non-Edge	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Auto
6	<input type="checkbox"/>	Auto	128	Non-Edge	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Auto

This page lets you view and/or edit the current STP CIST port configurations and contains settings for physical (Normal) and aggregated ports.

Point-to-Point and Edge Ports selections apply to RSTP only. Part of the task of configuring RSTP is defining the port types on the bridge, which is directly related to the device(s) connected to the port. With the port types defined, RSTP can reconfigure a network much quicker than STP when a change in network topology is detected. The two possible selections are **Non-Edge** (Point-to-point) port or **Edge** port.

If a bridge port is CIST connected to another bridge or router port, it normally operates in full-duplex mode and is functioning as a point-to-point port.

A port operates as an edge port when it is connected to a network terminal device such as a workstation or a server. An edge port on a bridge should not have any STP or RSTP devices connected to it either directly or through another device connected to that port. In this configuration since the port has no STP or RSTP devices connected to it, it will always forward network traffic.

Port

The LIB-44xx port number of the logical STP port (e.g., 1-4). The * in the Port column acts as a 'wild card' character which causes the selections in this row to be applied to all other Ports (rows) in the table for which this selection is valid. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

STP Enabled

Controls whether STP is enabled on this LIB-44xx port. Note that certain configurations require STP to be disabled.

Path Cost

Controls the path cost incurred by the port (dropdown and entry field).

The **Auto** setting will set the path cost as appropriate by the physical link speed, using the 802.1D recommended values. This is the default setting.

The **Specific** setting lets you define the Path Cost value in the entry box. 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 of 1 to 200,000,000. If **Specific** is selected at the dropdown, you must enter a Path Cost value in the entry field.

Priority

Controls the port priority. This can be used to control priority of ports having identical Path Costs (see above). Select 0, 16, 32, 48, 64, 80, 96, 112, 128, 144, 160, 176, 192, 208, 224, or 240. The default is **128**.

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 Admin Edge and Auto Edge 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). Select **Edge** or **Non-Edge**. The default is **Edge**.

AutoEdge

Controls whether the bridge should enable automatic edge detection on the bridge port. This allows *operEdge* to be derived from whether BPDUs are received on the port or not. The default is automatic edge detection on the bridge port enabled (checkbox checked).

Restricted Role

If checked, 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**. The default is unchecked.

Restricted TCN

If checked, 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 that the attached LANs transits frequently. The default is unchecked.

BPDU Guard

If checked, causes the port to disable itself upon receiving valid BPDUs. Contrary to the similar bridge setting, the port **Edge** status does not affect this setting. The default is unchecked.

A port entering error-disabled state due to this setting is subject to the bridge 'Port Error Recovery' setting as well. See the "Port Error Recovery" field description in the "[Bridge Settings](#)" section on page [120](#).

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.

Forced True: the port connects to a point-to-point LAN.

Forced False: the port connects to a shared medium.

Auto: the selection of whether the port connects to a point-to-point LAN or to a shared medium is automatically defined. The default is **Auto**.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

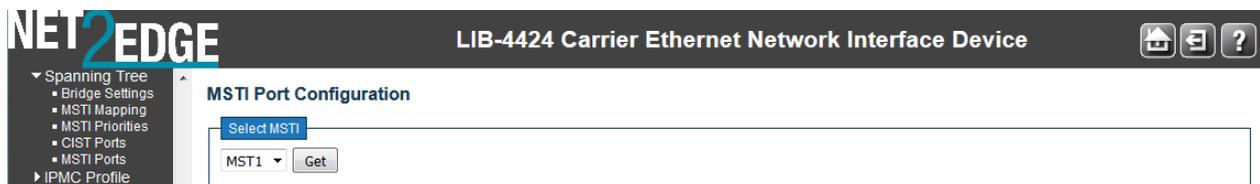
Note:

If you try to enable STP at the **Configuration > Spanning Tree > CST Ports** menu path while the 802.1X **Admin State** is set to any setting other than 'Force Authorized' at **Configuration > Security > Network > NAS**, the message "**STP Error - STP port configuration error**" displays.

You can set the 802.1X **Admin State** to a setting other than 'Force Authorized' at the **Configuration > Security > Network > NAS** menu path. See "[Configuration > Security > Network > NAS](#)" on page 65 for more information.

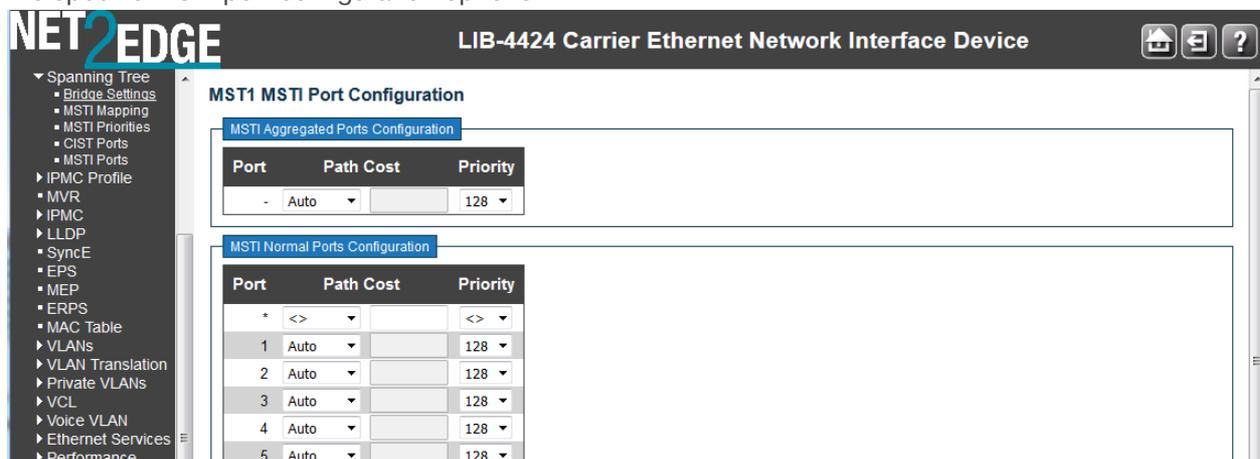
MSTI Ports

MSTI Port Configuration is done from the **Configuration > Spanning Tree > MSTI Ports** menu path.



This page lets you view and/or edit the current STP MSTI port configurations. 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.

Select the MSTI instance (MST1, MST2, etc.) from the dropdown and click the **Get** button to display the specific MSTI port configuration options:



This page contains MSTI port settings for physical and aggregated ports.

Port

The LIB-44xx port number of the corresponding STP CIST (and MSTI) port. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

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 **1** to **200,000,000**.

Priority

Controls the port priority. This can be used to control the priority of ports having identical port cost. (See above.) At the dropdown, select **0, 16, 32, 48, 64, 80, 96, 112, 128, 144, 160, 176, 192, 208, 224, or 240**.

The default is **128**.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

MVR Configuration

This page provides MVR related [configuration](#) from the **Configuration > MVR** menu path.

You can [view](#) Statistics, MVR Channel Groups, and MVR SFM Information from the **Monitor > MVR** menu path.

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 them.

The MVR feature enables multicast traffic forwarding on the Multicast VLANs. 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. You can create up to eight MVR VLANs with corresponding channel settings for each Multicast VLAN. There can be up to 256 group addresses for channel settings.

In multicast VLAN networks, subscribers to a multicast group can exist in multiple VLANs. If VLAN boundaries in a network consist of Layer 2 switches, it may be necessary to replicate the multicast stream to the same group in different subnets, even if they are on the same physical network. Multicast VLAN Registration (MVR) routes packets received in a multicast source VLAN to one or more receive VLANs. The clients are in the 'receive' VLANs and the multicast server is in the 'source' VLAN. Note that Multicast routing must be disabled when MVR is enabled.

IP Multicast delivers application source traffic to multiple receivers while minimizing burdening on the source / receivers. IP Multicast uses minimal network bandwidth and enables simple, scalable, economical applications distribution.

Multicast Virtual LAN Registration (MVR) increases multicast transport efficiency and is important for residential providers. MVR involves the creation of separate, dedicated VLANs constructed specifically for multicast traffic distribution. Each switch that receives an MVR stream examines each multicast group, and internally bridges the multicast VLAN traffic to a particular subscriber that has requested the specific multicast stream. This helps providers offer new, incremental services to their customers.

From the default page, click the **Add New MVR VLAN** button to display the edit fields.

VLAN Interface Setting (Role [I:Inactive / S:Source / R:Receiver])

MVR Mode

Enable or 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. The default is **Disabled**.

Delete

Check to delete the entry. The designated entry will be deleted during the next save.

MVR VID

Specify the Multicast VLAN ID (1-4094). **Caution:** MVR source ports should not be overlapped with Management VLAN ports.

MVR Name

Enter the name of the MVR VLAN. The maximum MVR VLAN Name string length is 32 characters. The MVR VLAN Name can only contain alpha or numeric characters, and it must contain at least one alpha character. The MVR VLAN Name can be edited for the existing MVR VLAN entries or it can be added to the new entries.

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 an MVR VID. The default is **Untagged**.

Priority

Specify how the traversed IGMP/MLD control frames will be sent in prioritized manner. The valid Priority range is **0** to **7**. The default Priority is **0**.

LLQI

Define 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 of a second (one-half second).

The Last Listener Query Interval] (LLQI) value specifies the maximum time allowed before sending a responding Report. Smaller LMQC/LLQC give smaller LMQT/LLQT; this condition shortens the leave latencies.

Interface Channel Setting

When the MVR VLAN is created, click the Edit symbol to expand the corresponding multicast channel settings for the specific MVR VLAN. A summary of the Interface Channel Setting (of the MVR VLAN) will display next to the Edit (e) symbol.

Immediate Leave Setting

Port

The logical port for the settings (e.g., ports 1-4 on the LIB-4400). Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Port Role

Configure an MVR port of the designated MVR VLAN as one of the following roles.

Inactive: The designated port does not participate in 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.

Select the port role by clicking the Role symbol to switch the setting.

I indicates Inactive (the default Role),

S indicates Source role, and

R indicates Receiver role.

VLAN Interface Setting (Role [I]:Inactive)

Delete	MVR VID	MVR
Delete		
Port	1 2 3 4 5	
Role	S R I I I	
Add New MVR VLAN		Inactive

Immediate Leave

Enable the fast leave function on the port.

Buttons

Add New MVR VLAN: Click to add new MVR VLAN. Specify the VID and configure the new entry. Edit the parameters as described above. Click **Save** when done.

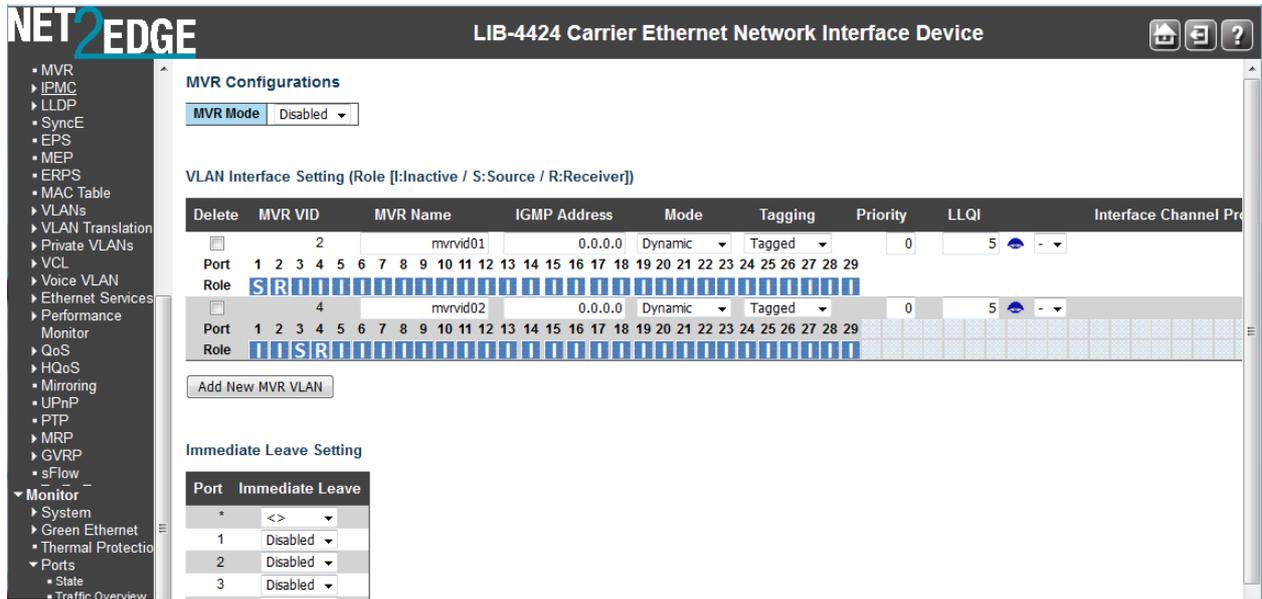
Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

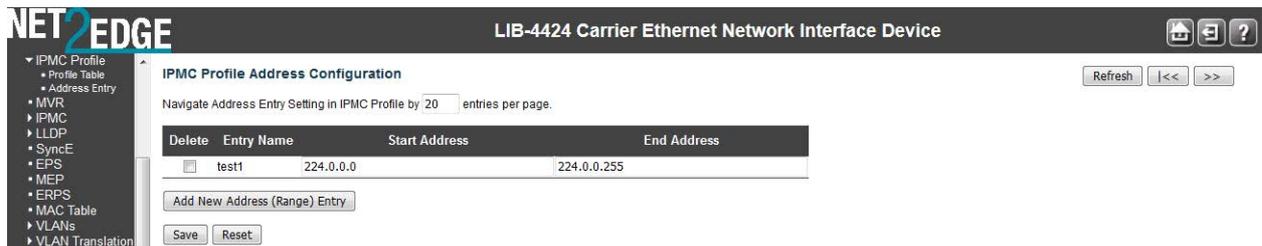
Edit (e): Click on the edit button to edit the related entry's channel configuration.

Example

The screen below shows two MVR VLANs configured (*mvrvid01* and *mvrvid02*).



MVR Interface Channel Profile using IPMC Profile. Access to IPMC Profile, Profile Table and Address Entry to add relevant data as below:



The IPMC Profile Address Configuration fields are explained below:

Delete

Click to delete the entry. The designated entry is deleted immediately.

Entry Name

Display the name of the specific Entry.

Start Address

The starting IPv4/IPv6 Multicast Group Address to be used as a streaming channel. For example:

<u>Start</u>	<u>End</u>	<u>Description</u>
224.0.0.0	224.0.0.255	Reserved for special “well-known” multicast addresses.
224.0.1.0	238.255.255.255	Globally-scoped (Internet-wide) multicast addresses.
239.0.0.0	239.255.255.255	Administratively-scoped (local) multicast addresses.

See the IANA web site at <http://www.iana.org/assignments/multicast-addresses/multicast-addresses.xml> or <http://www.iana.org/assignments/ipv6-multicast-addresses/ipv6-multicast-addresses.xml>

[addresses.xml](#) for the full set of address assignments. IPv6 multicast addresses are defined in "IP Version 6 Addressing Architecture" per IETF [\[RFC4291\]](#).

End Address

The ending IPv4/IPv6 Multicast Group Address to be used as a streaming channel.

IP multicast is a way of sending Internet Protocol (IP) datagrams to a group of interested receivers in a single transmission. It is often employed for streaming media applications on the Internet and private networks. The method is the IP-specific version of the general concept of multicast networking. It uses specially reserved multicast address blocks in IPv4 and IPv6. In IPv6, IP multicast addressing replaces broadcast addressing as implemented in IPv4.

Traditional IP communication lets a host send packets to a single host (unicast transmission) or to all hosts (broadcast transmission). IPv6 multicast provides a third scheme, allowing a host to send a single datastream to a subset of all hosts (group transmission) simultaneously.

This page provides IPMC Profile Rule Setting

IPMC Profile [test] Rule Settings (In Precedence Order)

Profile Name & Index	Entry Name	Address Range	Action	Log
test 1	test1	224.0.0.0 ~ 224.0.0.255	Permit	Disable

MVR Configurations

MVR Mode: Enabled

VLAN Interface Setting (Role [I:Inactive / S:Source / R:Receiver])

Delete	MVR VID	MVR Name	IGMP Address	Mode	Tagging	Priority	LLQI	Interface Channel Profile
<input type="checkbox"/>	2	mvrvid01	0.0.0.0	Dynamic	Tagged	0	5	test
<input type="checkbox"/>	4	mvrvid02	0.0.0.0	Dynamic	Tagged	0	5	-

Immediate Leave Setting

Port	Immediate Leave
*	<>
1	Enabled
2	Enabled
3	Disabled

IPMC (IP MultiCast)

IP Multicast (IPMC) is a way to send Internet Protocol (IP) datagrams to a group of interested receivers in a single transmission. It is often used for streaming media applications on the Internet and private networks. The method is the IP-specific version of the general concept of multicast networking. It uses specially reserved multicast address blocks in IPv4 and IPv6. In IPv6, IP multicast addressing replaces broadcast addressing as implemented in IPv4. IP multicast is used in enterprises, commercial stock exchanges, and multimedia content delivery networks. One common enterprise use of IP multicast is for IPTV applications such as distance learning and televised company meetings. Multicast is a different transmission mode from unicast, so only protocols designed for multicast are used with multicast.

The LIB-44xx IPMC menu provides IGMP Snooping and MLD Snooping configuration from the **Configuration > IPMC** menu path. These two sub-menus are described in the following sections. The Internet Group Management Protocol (IGMP) communications protocol is used to manage the membership of Internet Protocol multicast groups. IGMP is used by IP hosts and adjacent multicast routers to establish multicast group memberships. It is an integral part of the IP multicast specification, like ICMP for unicast connections. IGMP allows more efficient use of resources when supporting online video, gaming, etc.

The LIB-44xx can do IGMPv1/v2/v3 and MLDv1/v2 snooping to limit the broadcast of the IGMP multicast sessions to the ports where the IGMP listeners can be reached. MLD is similar to IGMP except MLD runs over the IPv6 stack. The LIB-44xx looks for IGMP 'join' and 'leave' messages and maintains a table of which ports are part of the conversation. Snooping is enabled at device level and also supports proxying. This feature can be used to avoid forwarding unnecessary join and leave messages to the router interface.

MLD Snooping Configuration

Global Configuration	
Snooping Enabled	<input type="checkbox"/>
Unregistered IPMCv6 Flooding Enabled	<input checked="" type="checkbox"/>
MLD SSM Range	ff3e:: / 96
Leave Proxy Enabled	<input type="checkbox"/>
Proxy Enabled	<input type="checkbox"/>

Port Related Configuration

Port	Router Port	Fast Leave	Throttling
*	<input type="checkbox"/>	<input type="checkbox"/>	<>
1	<input type="checkbox"/>	<input type="checkbox"/>	unlimited
2	<input type="checkbox"/>	<input type="checkbox"/>	unlimited
3	<input type="checkbox"/>	<input type="checkbox"/>	unlimited
4	<input type="checkbox"/>	<input type="checkbox"/>	unlimited

IGMP Snooping Configuration

Global Configuration	
Snooping Enabled	<input type="checkbox"/>
Unregistered IPMCv4 Flooding Enabled	<input checked="" type="checkbox"/>
IGMP SSM Range	232.0.0.0 / 8
Leave Proxy Enabled	<input type="checkbox"/>
Proxy Enabled	<input type="checkbox"/>

Port Related Configuration

Port	Router Port	Fast Leave	Throttling
*	<input type="checkbox"/>	<input type="checkbox"/>	<>
1	<input type="checkbox"/>	<input type="checkbox"/>	unlimited
2	<input type="checkbox"/>	<input type="checkbox"/>	unlimited
3	<input type="checkbox"/>	<input type="checkbox"/>	unlimited
4	<input type="checkbox"/>	<input type="checkbox"/>	unlimited

The LIB-44xx also supports IGMP/MLD snooping at VLAN levels. A maximum of 64 VLANs can be chosen for IGMP snooping. Typically the router is the IGMP querier but an option to enable IGMP querier on each VLAN is provided as well on this device. The IGMP querier will send a query in 255 seconds after enabled; if it receives any query from other devices, this will stop querying.

MLD Snooping VLAN Configuration

Start from VLAN with entries per page.

VLAN ID	Snooping Enabled	MLD Querier
1	<input checked="" type="checkbox"/>	<input type="checkbox"/>
100	<input type="checkbox"/>	<input type="checkbox"/>

IGMP Snooping VLAN Configuration

Start from VLAN with entries per page.

VLAN ID	Snooping Enabled	IGMP Querier
1	<input type="checkbox"/>	<input type="checkbox"/>
100	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

The LIB-44xx provides status on the IGMP sessions and statistics of different queries and messages as discussed later in this section.

IGMP Snooping

The IGMP Snooping menu provides for Basic Configuration, VLAN Configuration, and Port Group Filtering configuration from the sub-menus.

IGMP snooping allows the S3240 dynamically determine which hosts connected to a particular VLAN in the switch need to receive a particular multicast transmission. The S3240 basically listens (snoops) to the various IGMP messages (e.g., 'Query' or 'Leave') and other multicast protocol transmissions. It then dynamically determines which egress ports are associated with each multicast transmission. The LIB-44xx uses a bridge table entry to control multicast forwarding (note that the entry is dynamically configured). The LIB-44xx performs these actions based on IGMP messages snooped: add a receiver to a group, remove a receiver from a group, or maintain group membership.

Basic Configuration

From the **Configuration > IPMC > IGMP Snooping > Basic Configuration** menu path you can view and edit the IGMP Snooping global and port-related configurations.

Global Configuration			
Snooping Enabled	<input checked="" type="checkbox"/>		
Unregistered IPMCv4 Flooding Enabled	<input checked="" type="checkbox"/>		
IGMP SSM Range	232.0.0.0	/ 8	
Leave Proxy Enabled	<input type="checkbox"/>		
Proxy Enabled	<input type="checkbox"/>		

Port Related Configuration			
Port	Router Port	Fast Leave	Throttling
*	<input type="checkbox"/>	<input type="checkbox"/>	<>
1	<input type="checkbox"/>	<input checked="" type="checkbox"/>	unlimited
2	<input type="checkbox"/>	<input checked="" type="checkbox"/>	unlimited
3	<input type="checkbox"/>	<input checked="" type="checkbox"/>	unlimited
4	<input type="checkbox"/>	<input checked="" type="checkbox"/>	unlimited
5	<input type="checkbox"/>	<input checked="" type="checkbox"/>	unlimited
6	<input type="checkbox"/>	<input checked="" type="checkbox"/>	unlimited

This page provides IGMP Global and IGMP Port Related Snooping configuration.

Snooping Enabled

Check to enable Global IGMP Snooping. The default is unchecked (disabled).

Unregistered IPMCv4 Flooding Enabled

Check to enable unregistered IPMC traffic flooding. The default is enabled (checkbox checked). IPMC IP MultiCast (IPMC) supports IPv4 and IPv6 multicasting. IPMCv4 denotes multicast for IPv4. IPMCv6 denotes multicast for IPv6.

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. SSM is a method of delivering multicast packets in which the only packets delivered to a receiver are those originating from a specific source address requested by the receiver. By so limiting the source, SSM reduces demands on the network and improves security. SSM requires that the receiver specify the source address and explicitly excludes the use of the (*,G) join for all multicast groups in [RFC 3376](#), which is possible only in IPv4's IGMPv3 and IPv6's MLDv2.

SSM can be viewed in contrast to ASM (Any-Source Multicast), where a receiver expresses interest in traffic to a multicast address. The multicast network must 1) discover all multicast sources sending to that address, and 2) route data from all sources to all interested receivers. ASM is particularly well suited to groupware applications where 1) all participants in the group want to be aware of all other participants, and 2) the list of participants is not known in advance. With ASM, the source discovery burden on the network can become significant with a large number of sources.

With SSM, the receiver expresses interest in traffic to a multicast address, and also expresses interest in receiving traffic from just one specific source sending to that multicast address. This keeps the network from having to discover numerous multicast sources, and reduces the amount of multicast routing information that the network must maintain. SSM requires support in last-hop routers and in the receiver's operating system. SSM support is not required in other network components (including routers and even the sending host). Interest in multicast traffic from a specific source is conveyed from hosts to routers using IGMPv3 as specified in [RFC 4607](#). In SSM, some types of denial of service (DoS) attack cannot be made by simply sending traffic to a multicast group.

SSM destination addresses must be in the ranges 232.0.0.0/8 for IPv4 or FF3x::/96 for IPv6. See <http://tools.ietf.org/html/rfc4607> for the full set of reserved addresses.

Leave Proxy Enabled

Check to enable IGMP Leave Proxy. This feature can be used to avoid forwarding unnecessary leave messages to the router side. The default is unchecked.

Proxy Enabled

Check to enable IGMP Proxy. This feature can be used to avoid forwarding unnecessary join and leave messages to the router side. The default is unchecked.

Port

The LIB-44xx logical port number (e.g., 1-4 on the LIB-4400). The * in the Port column acts as a 'wild card' character which causes the selections in this row to be applied to all other Ports (rows) in the table for which this selection is valid. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Router Port

Specify which ports act as router ports. A router port is an Ethernet port on the LIB-44xx 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. The default is unchecked.

Fast Leave

Check to enable the fast leave on the port. With IGMP fast-leave processing enabled, the LIB-44xx immediately removes the interface attached to a receiver on reception of a Leave Group message. This speeds up leave processing, but should only be used when receivers are directly attached to the LIB-44xx. The default is unchecked (fast leave disabled on the port).

When you enable IGMP fast-leave processing, the LIB-44xx immediately removes a port when it detects an IGMP v2 leave message on that port.

Throttling

Select **unlimited** or a value from **1 - 10** to limit the number of multicast groups to which an LIB-44xx port can belong. The default is **unlimited**.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

VLAN Configuration

The **Configuration > IPMC > IGMP Snooping > VLAN Configuration** menu path lets you view and edit the IGMP Snooping VLAN Configuration table. From the default page, click the Add New IGMP VLAN button to display the edit fields.

Each page shows up to 99 entries from the VLAN table (default of 20) selected through the "entries per page" input field. When first visited, the page shows 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.

The "VLAN" input fields let you select the starting point in the VLAN Table. Clicking the **Refresh** button will update the displayed table starting from that or the next closest VLAN Table match. The **>>** button will use the last entry of the currently displayed entry as a basis for the next lookup. When the end is reached, the text "*No more entries*" displays in the table. Use the **<<** button to start over.

IGMP Snooping VLAN Table Columns

VLAN ID

The VLAN ID of the entry.

IGMP Snooping Enabled

Check to enable the per-VLAN IGMP Snooping. Up to 64 VLANs can be selected. The default is unchecked (per-VLAN IGMP Snooping disabled).

IGMP Querier

Check to enable the IGMP Querier in the VLAN. The IGMP 'Querier' is a router that sends IGMP Query messages onto a particular link. The default is unchecked (IGMP Querier disabled). The IGMP snooping querier supports IGMP Versions 1 and 2. The IGMP snooping querier becomes disabled if it detects a multicast router in the network.

Compatibility

Select **IGMP-Auto**, **Forced IGMPv1**, **Forced IGMPv2**, or **Forced IGMPv3**. 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 default compatibility value is **IGMP-Auto**.

IGMP-Auto: Compatibility is automatically assigned.

Forced IGMPv1: Compatibility is forced to IGMP version 1.

Forced IGMPv2: Compatibility is forced to IGMP version 2.

Forced IGMPv3: Compatibility is forced to IGMP version 3.

Three versions of IGMP exist - versions v1, v2, and v3. One difference between the versions is how a host node signals that it no longer wants to be a member of a multicast group. In IGMP v1, the host node stops sending reports. If a router does not receive a report from a host node after a predefined length of time (time-out value) it assumes that the host node no longer wants to receive multicast frames and removes it from the membership list of the multicast group. In version 2, a host node exits from a multicast group by sending a leave request. After receiving a leave request from a host node, the router removes the node from appropriate membership list. The router also stops sending multicast packets from the port if it determines there are no further host nodes on the port. Version 3 adds the ability of host nodes to "join" or "leave" specific sources in a multicast group.

RV

Displays the Robustness Variable (RV) which allows tuning for the expected packet loss on a network.

The valid range is **1** to **255**. The default robustness variable value is **2**.

QI (sec)

Displays the Query Interval. The Query Interval (QI) is the interval between General Queries sent by the Querier. The valid range is **1** to **255** seconds. The default query interval is **125** seconds.

QRI (0.1 sec)

Displays the Query Response Interval. The Max Response Time used to calculate the Max Resp Code inserted into the periodic General Queries. The valid range is **0** to **31744** in tenths of a second. The default query response interval is **100** in tenths of a second (10 seconds).

LLQI (0.1 sec) (LMQI for IGMP)

Displays the 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 valid range is **0** to **31744** in tenths of a second. The default last member query interval is **10** in tenths of a second (1 second).

URI (sec)

Displays the Unsolicited Report Interval. The Unsolicited Report Interval (URI) is the time between repetitions of a host's initial report of membership in a group. The valid range is **0** to **31744** seconds. The default unsolicited report interval is **1** second.

Buttons

Refresh: Refreshes the displayed table starting from the "VLAN" input fields.

<<: Updates the table starting at the first VLAN Table entry (the entry with the lowest VLAN ID).

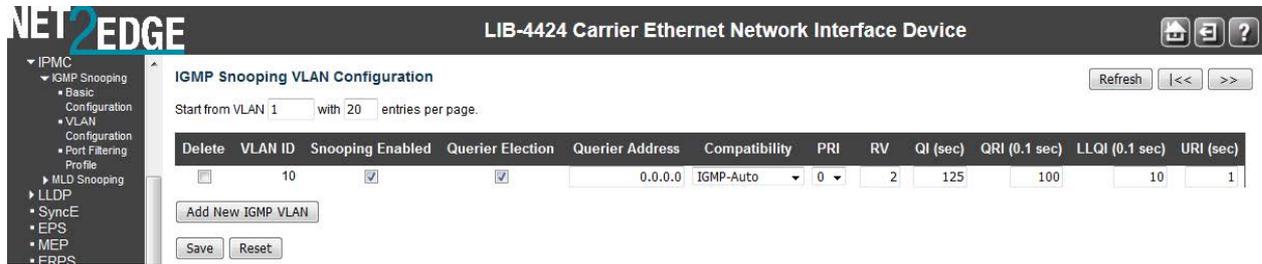
>>: Updates the table, starting with the entry after the last entry currently displayed.

Save: Click to save changes (required to display the

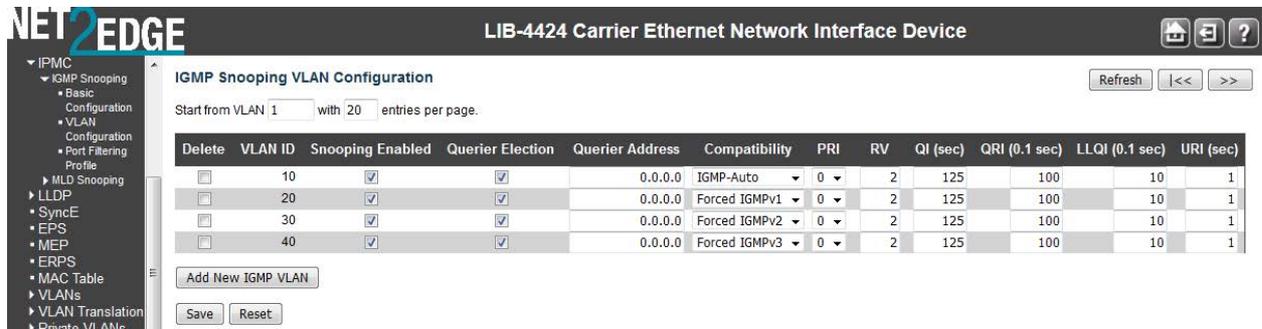
Reset: Click to undo any changes made locally and revert to previously saved values.

Example

A sample **Configuration > IPMC > IGMP Snooping > VLAN Configuration** is shown below: In this example, VLAN ID 1 has Snooping and IGMP Querier enabled, and has Forced IGMPv1 compatibility selected.



In the example below, VLAN IDs 1, 10 and 11 have Snooping enabled with various Compatibility settings (Forced IGMPv1, IGMP-Auto, and Forced IGMPv3). VLAN IDs 1, 10, 11, 12 and 13 have IGMP Querier enabled, with default settings for RV, QI, QRI, LLQI, and URI.

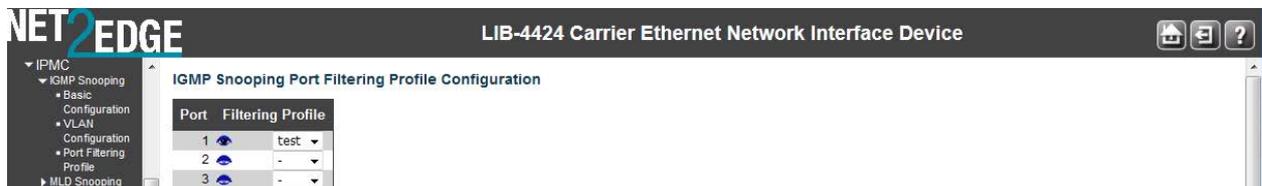


Note that VLAN IDs 12 and 13 have Snooping Enabled set to “disabled”.

Port Group Filtering

From the **Configuration > IPMC > IGMP Snooping > Port Group Filtering Configuration** menu path you can view and edit the IGMP Snooping Port Group Filtering Configuration table.

From the default page, click the **Add New Filtering Group** button to display the edit fields.

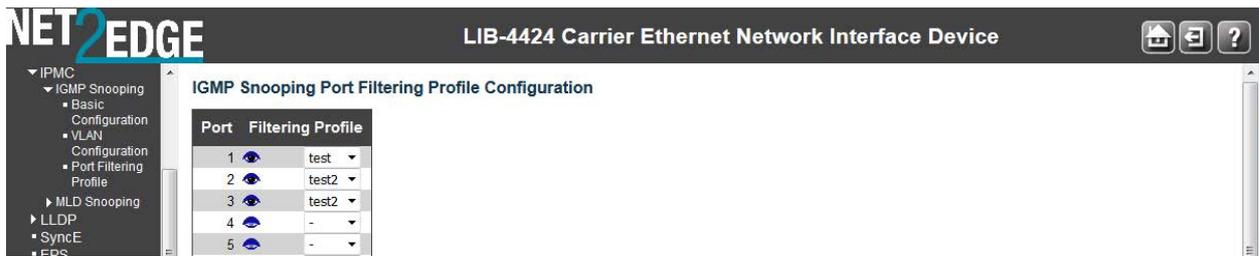


Port

The logical port for the settings. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Filtering Profile

The IPMC Profile Table entry



Note in the example above that a Port cannot have the same Filtering Group assigned twice. If you try to do so, the entry is not accepted after you click 'Save'. Note however that two different Ports can have the same Filtering Group; both Port 1 and Port 2 have Filtering Group 225.168.11.11 assigned in the sample screen above.

Click on , you can view the configured Filtering Multicast Address with relevant parameters.

IPMC Profile [test2] Rule Settings (In Precedence Order)

Profile Name & Index	Entry Name	Address Range	Action	Log
test2	1 test2	224.1.0.0 ~ 224.1.0.255	Permit	Disable

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Add New Filtering Group: Click to add a new entry to the Group Filtering table.

MLD Snooping

The MLD Snooping menu provides for Basic Configuration, VLAN Configuration, and Port Group Filtering configuration from the sub-menus from the **Configuration > IPMC > MLD Snooping** menu path.

Multicast Listener Discovery 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. MLD snooping is an IPv6 multicast constraining mechanism that runs on Layer 2 devices to manage and control IPv6 multicast groups. By analysing received MLD messages, a Layer 2 device running MLD snooping establishes mappings between ports and multicast MAC addresses and forwards IPv6 multicast data based on these mappings. When MLD snooping is not running, IPv6 multicast packets are broadcast to all devices on Layer 2. With MLD snooping running, multicast packets for known IPv6 multicast groups are multicast to the receivers at Layer 2.

MLD snooping forwards multicast data to only the receivers requiring it at Layer 2, providing advantages such as reducing Layer 2 broadcast packets for network bandwidth savings, enhancing multicast traffic security, and providing per-host accounting.

Basic Configuration

From the **Configuration > IPMC > MLD Snooping > Basic Configuration** menu path you can view and edit the MLD Snooping global and port configurations.

MLD Snooping Configuration

Global Configuration

Snooping Enabled	<input checked="" type="checkbox"/>
Unregistered IPMCv6 Flooding Enabled	<input checked="" type="checkbox"/>
MLD SSM Range	ff3e:: / 96
Leave Proxy Enabled	<input type="checkbox"/>
Proxy Enabled	<input type="checkbox"/>

Port Related Configuration

Port	Router Port	Fast Leave	Throttling
*	<input type="checkbox"/>	<input type="checkbox"/>	<>
1	<input type="checkbox"/>	<input type="checkbox"/>	unlimited
2	<input type="checkbox"/>	<input type="checkbox"/>	unlimited
3	<input type="checkbox"/>	<input type="checkbox"/>	unlimited
4	<input type="checkbox"/>	<input type="checkbox"/>	unlimited
5	<input type="checkbox"/>	<input type="checkbox"/>	unlimited

This page provides MLD Snooping related configuration at the global level and at the port level.

Snooping Enabled

Check to enable Global MLD Snooping. The default is unchecked (snooping disabled).

Unregistered IPMC Flooding Enabled

Enable unregistered IPMCv6 traffic flooding. **Note:** disabling unregistered IPMCv6 traffic flooding may cause Neighbour Discovery failure. The default is checked (enabled).

MLD SSM Range

SSM (Source-Specific Multicast) Range allows the SSM-aware hosts and routers to run the SSM service model for the groups in the address range. SSM destination addresses must be in the ranges 232.0.0.0/8 for IPv4 or FF3x::/96 for IPv6.

SSM Range allows the SSM-aware hosts and routers run the SSM service model for the groups in the address range. SSM is a method of delivering multicast packets in which the only packets delivered to a receiver are those originating from a specific source address requested by the receiver. By so limiting the source, SSM reduces demands on the network and improves security.

SSM requires that the receiver specify the source address and explicitly excludes the use of the (*,G) join for all multicast groups in [RFC 3376](#), which is possible only in IPv4's IGMPv3 and IPv6's MLDv2.

SSM is best viewed in contrast to ASM (Any-Source Multicast), where a receiver expresses interest in traffic to a multicast address. The multicast network must 1) discover all multicast sources sending to that address, and 2) route data from all sources to all interested receivers. ASM is particularly well suited to groupware applications where 1) all participants in the group want to be aware of all other participants, and 2) the list of participants is not known in advance. With ASM, the source discovery burden on the network can become significant with a large number of sources.

With SSM, the receiver expresses interest in traffic to a multicast address, and also expresses interest in receiving traffic from just one specific source sending to that multicast address. This keeps the network from having to discover numerous multicast sources, and reduces the amount of multicast routing information that the network must maintain. SSM requires support in last-hop routers and in the receiver's operating system. SSM support is not required in other network components (including routers and even the sending host). Interest in multicast traffic from a specific source is conveyed from hosts to routers using IGMPv3 as specified in [RFC 4607](#). In SSM, some

types of denial of service (DoS) attack cannot be made by simply sending traffic to a multicast group.

SSM destination addresses must be in the ranges 232.0.0.0/8 for IPv4 or FF3x::/96 for IPv6. See <http://tools.ietf.org/html/rfc4607> for the full set of reserved addresses.

Leave Proxy Enabled

Check to enable MLD Leave Proxy. This feature can be used to avoid forwarding unnecessary Leave messages to the router side. The default is unchecked.

Proxy Enabled

Check to enable MLD Proxy. This feature can be used to avoid forwarding unnecessary Join and Leave messages to the router side. The default is unchecked.

Port

The LIB-44xx logical port number (e.g., 1-4 on the LIB-4400). The * in the Port column acts as a 'wild card' character which causes the selections in this row to be applied to all other Ports (rows) in the table for which this selection is valid. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Router Port

Check to specify which ports act as router ports. A router port is a port on the LIB-44xx 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. The default is unchecked.

Fast Leave

Check to enable the fast leave function on the related port. The default is unchecked. Multicast snooping Fast Leave processing allows the LIB-44xx to remove an interface from the forwarding table entry without first sending out group specific queries to the interface. The VLAN interface is 'pruned' from the multicast tree for the multicast group specified in the original leave message. Fast leave is enabled at the port level. Pruning happens per VLAN per port. See the IGMP group table which is indexed by VLAN and group. Fast-leave processing ensures optimal bandwidth management for all hosts on a switched network, even when multiple multicast groups are in use simultaneously. Fast Leave processing applies to both IGMP and MLD. When you enable MLD fast-leave processing, the LIB-44xx immediately removes a port when it detects an IGMP v2 leave message on that port.

Throttling

Used to limit the number of multicast groups to which a LIB-44xx port can belong. Select **unlimited** or **1-10** multicast groups as the limit. The default is **unlimited**.

Buttons

Save: Click to save changes (required to be able to edit all fields).

Reset: Click to undo any changes made locally and revert to previously saved values.

MLD Snooping > VLAN Configuration

From the **Configuration > IPMC > MLD Snooping > VLAN Configuration** menu path you can view and edit the MLD Snooping VLAN Configuration parameters. At the default page, click the Add New MLD VLAN button to display the edit fields.

Each page shows up to 99 entries from the VLAN table (default of 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.

The "VLAN" input fields allow the user to select the starting point in the VLAN Table. Clicking the **Refresh** button will update the displayed table starting from that or the next closest VLAN Table match.

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

The MLD Snooping VLAN table columns are explained below:

VLAN ID

The VLAN ID of the entry.

Snooping Enabled

Check to enable the per-VLAN MLD Snooping. Up to 64 VLANs can be selected. The default is unchecked.

MLD Querier

Check to enable the IGMP Querier in the VLAN. The default is unchecked.

Compatibility

Select **MLD-Auto**, **Forced MLDv1**, or **Forced MLDv2**. 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 default compatibility value is **MLD-Auto**.

MLD-Auto: Compatibility is automatically assigned.

Forced MLDv1: Compatibility is forced to MLD version 1. MLD v1 was the original release of MLD as an asymmetric protocol, specifying different behaviours for multicast listeners and for routers per IETF [RFC 2710](#).

Forced MLDv2: Compatibility is forced to MLD version 2. MLDv2 is designed to be interoperable with MLDv1. MLDv2 adds the ability for a node to report interest in listening to packets with a particular multicast address only from specific source addresses or from all sources except for specific source addresses. Refer to IETF [RFC 3810](#).

RV

The Robustness Variable allows tuning for the expected packet loss on a link. The valid range is **1** to **255**, default robustness variable value is **2**.

QI (sec.)

The Query Interval variable - denotes the interval between General Queries sent by the Querier. The valid range is **1** to **255** seconds. The default query interval is **125** seconds.

QRI (0.1 sec.)

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 a second. The default query response interval is **100** in tenths of a second (10 seconds).

LLQI (0.1 sec.)

The Last Listener Query Interval - 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 valid range is **0** to **31744** in tenths of a second. The default last listener query interval is **10** in tenths of a second (1 second).

URI (sec.)

The Unsolicited Report Interval - the time between repetitions of a node's initial report of interest in a multicast address. The valid range is **0** to **31744** seconds. The default URI is **1** second.

Buttons

Refresh: Refreshes the displayed table starting from the "VLAN" input fields.

<<: Updates the table starting from the first entry in the VLAN Table (i.e., the entry with the lowest VLAN ID).

>>: Updates the table, starting with the entry after the last entry currently displayed.

Example

A sample **Configuration > IPMC > MLD Snooping > VLAN Configuration** is shown below: In this example, VLAN ID 1 has Snooping enabled and IGMP Querier disabled; VLAN ID 10 has Snooping disabled and IGMP Querier enabled.

Delete	VLAN ID	Snooping Enabled	Querier Election	Compatibility	PRI	RV	QI (sec)	QRI (0.1 sec)	LLQI (0.1 sec)	URI (sec)
<input type="checkbox"/>	10	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	MLD-Auto	0	2	125	100	10	1
<input type="checkbox"/>	11	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	MLD-Auto	0	2	125	100	10	1

Note that VLAN ID 10 has Snooping disabled in the example screen above.

Port Group Filtering

From the **Configuration > IPMC > MLD Snooping > Port Group Filtering** menu path you can add, view, and edit the MLD Snooping Port Group Filtering parameters.

At the default page, click the “**Add New Filtering Group**” button to display the editable table.



Port

Select the logical port for the settings (e.g., ports 1-4 on the LIB-4400). Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Filtering Profile

Select the available Filtering Profile, this selection is taken from IPMC Profile Entry.

In the example below, Port 1 has two different Filtering Groups assigned, and Port 2 has one Filtering Group assigned. Note that Port 1 and Port 2 cannot have the same address assigned for their Filtering Groups.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Add New Filtering Group: Click to add a new entry to the Group Filtering table.

Messages: *The same entry [Port 1, Group ff02::1:ff00:100] already exists.*

LLDP Configuration

The LIB-44xx **Configuration > LLDP** menu path lets you configure LLDP at the device level and at the port level.

The Link Layer Discovery Protocol (LLDP) IEEE 802.1ab standard protocol 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 LLDP 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).

LLDP information is sent as an Ethernet frame by devices from each of their interfaces at a fixed interval. Each frame contains one Link Layer Discovery Protocol Data Unit (LLDPDU). Each LLDPDU is a sequence of type-length-value (TLV) structures. The LLDP Ethernet frame typically has its destination MAC address set to a special multicast address that 802.1D-compliant bridges do not

forward (other multicast and unicast destination addresses are permitted). The EtherType field is set to 0x88cc.

Each LLDP frame starts with mandatory TLVs (Chassis ID, Port ID, and Time-to-Live). The mandatory TLVs are followed by a series of optional TLVs. The frame ends with the 'end of LLDPDU' with both its type and length fields set to 0.

The LLDP Configuration menu is available from the **Configuration > LLDP > LLDP** menu path.

NET2EDGE LIB-4424 Carrier Ethernet Network Interface Device

LLDP Configuration

LLDP Parameters

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

LLDP Interface Configuration

Interface	Mode	Optional TLVs						
		CDP aware	Trap	Port Descr	Sys Name	Sys Descr	Sys Capa	Mgmt Addr
*	<>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>				
GigabitEthernet 1/1	Enabled	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>				
GigabitEthernet 1/2	Enabled	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>				
GigabitEthernet 1/3	Enabled	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>				
GigabitEthernet 1/4	Enabled	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>				
GigabitEthernet 1/5	Enabled	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>				
GigabitEthernet 1/6	Enabled	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>				

This page lets you view and configure the LLDP parameters and port settings as explained below:

LLDP Parameters

Tx Interval

The LIB-44xx periodically transmits LLDP frames to its neighbours 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 **5** - **32768** seconds. The default is **30** seconds.

Tx Hold

Each LLDP frame contains information about how long the information in the LLDP frame are considered valid. The LLDP information valid period is set to **Tx Hold** multiplied by **Tx Interval** seconds. Valid values are **2** - **10** times. The default is **3** 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 **1** - **8192** seconds. The default is **2** seconds.

Tx Reinit

When a port is disabled, LLDP is disabled or the LIB-44xx is rebooted, an LLDP shutdown frame is transmitted to the neighbouring units, signalling that the LLDP information isn't valid anymore. **Tx Reinit** controls the amount of seconds between the shutdown frame and a new LLDP initialization. Valid values are **1** - **10** seconds. The default is **2** seconds.

LLDP Port Configuration

Port

The LIB-44xx port number of the logical LLDP port (e.g., 1-4 on the LIB-4400). The * in the Port column acts as a 'wild card' character which causes the selections in this row to be applied to all other Ports (rows) in the table for which this selection is valid. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Mode

Select LLDP mode. The valid selections are:

Disabled The LIB-44xx will not send out LLDP information, and will drop LLDP information received from neighbours.

Enabled The LIB-44xx will send out LLDP information, and will analyse LLDP information received from neighbours.

Rx only The LIB-44xx will not send out LLDP information, but LLDP information from neighbour units is analysed.

Tx only The LIB-44xx will drop LLDP information received from neighbours, but will send out LLDP information.

CDP aware

Enable or disable CDP (Cisco Discovery Protocol) awareness. The default is disabled (checkbox unchecked). The CDP operation is restricted to decoding incoming CDP frames (The LIB-44xx doesn't transmit CDP frames). CDP frames are only decoded if LLDP on the port is enabled.

Only CDP TLVs that can be mapped to a corresponding field in the LLDP neighbours' 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 neighbours' 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 neighbours 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 neighbours' table.

If all ports have CDP awareness disabled the switch forwards CDP frames received from neighbour devices. If at least one port has CDP awareness enabled all CDP frames are terminated by the switch.

Note: When CDP awareness on a port is disabled, the CDP information is not 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.

An LLDP frame contains multiple TLVs. 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 TLV is disabled, the corresponding information is not included in the LLDP frame.

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.

Buttons

Save: Click to save changes.

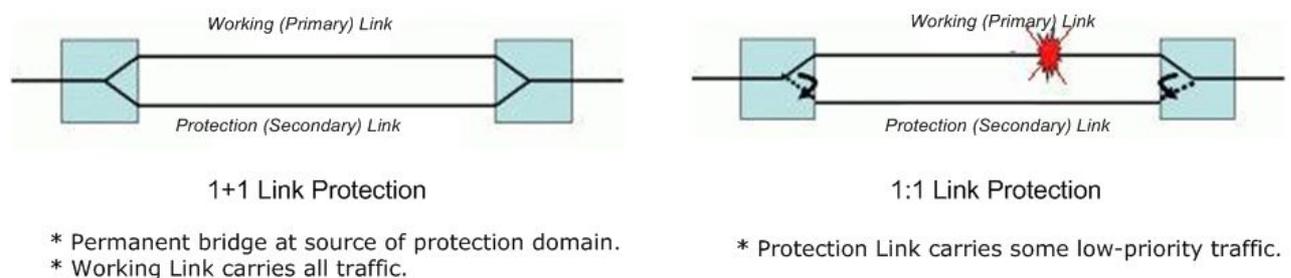
Reset: Click to undo any changes made locally and revert to previously saved values.

The LIB-44xx maintains LLDP port and protocol statistics and a table of neighbours that were discovered. See **Monitor > LLDP** for more information.

EPS Configuration

Ethernet (Linear) Protection Switch (EPS) instances are configured from the **Configuration > EPS** menu path.

The LIB-44xx implements the ITU G.8031 standard for EPS at the port level. It can perform 1:1 and 1+1 protection in unidirectional and bidirectional switching. The EPS feature provides the option to configure revertive and non-revertive mode in both 1:1 and 1+1 switching. It also allows configuring WTR (Wait To Restore) timer in revertive mode to avoid flapping working link which could trigger constant protection switching. The LIB-44xx supports a 'Hold-Off-Timer' which will delay the protection switching until an upstream device or the lower layer is ready. Both the WTR and Hold-Off-Timer are configurable in fine granular time increments.

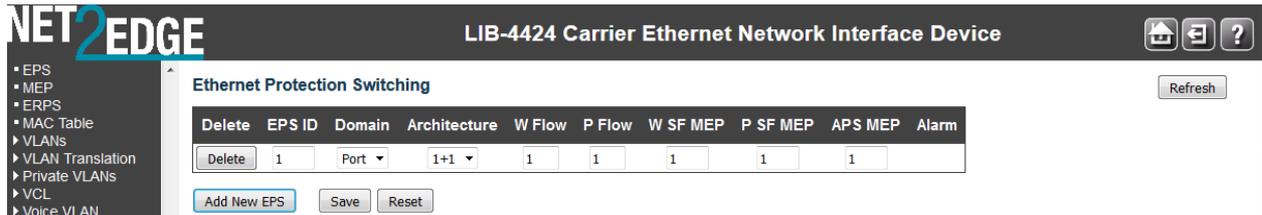


The EPS instance is associated with the MEPs on the working and protection links which are responsible for sending/receiving of APS frames. The MEPs associated with the EPS must have APS enabled for protection. APS frames can be unicast to the peer or a multicast. The APS frames can

carry specific commands to its peer entity. When there is a failure in the working link in 1+1 or 1:1 switching, the protection link takes over.

Note: Since the protection switching mechanism requires monitoring for both working and protection transport entities, MEPs must be activated for monitoring the working and protection transport entities. See the “[SOAM MEP Configuration](#)” section that follows this section for more information.

From the default page, click the **Add New EPS** button to display the Ethernet Protection Switching table.



The Ethernet (Linear) Protection Switch parameters are explained below:

Delete

This checkbox is used to mark an EPS for deletion in the next Save operation.

EPS ID

The ID of the EPS. Click on the ID of an EPS to enter its configuration page (see below).

Domain

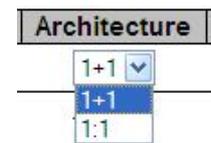
Port: This will create an EPS in the Port Domain. 'W Flow' and 'P Flow' are Ports.

Architecture

Select the linear protection switching architecture; either **1+1** protection switching or **1:1** protection switching architecture.

1+1: This will create a 1+1 EPS. The linear 1+1 protection switching architecture operates with either unidirectional or bidirectional switching. In a 1+1 architecture, a protection transport entity is used to protect the normal traffic signal. At the head-end, the bridge is permanent. Switching occurs only at the tail-end. 1+1 protection is often provisioned as non-revertive operation.

1:1: This will create a 1:1 (bidirectional) EPS. The linear 1:1 protection switching architecture operates with bidirectional switching. In a 1:1 architecture, a protection transport entity is used to protect the normal traffic signal. At the head-end, the bridge is not established until a protection switch is required. 1:1 protection is usually provisioned as revertive operation.



Note: The architecture at each end of the protected domain must match. Bidirectional switching always requires APS communication.

W Flow

Enter the working flow (W Flow) port for the EPS (1-4). See the 'Domain' parameter description above.

P Flow

Enter the protecting flow (P Flow) port for the EPS (1-4). See the 'Domain' parameter description above.

W SF MEP

Enter the working Signal Fail (SF) reporting MEP.

P SF MEP

Enter the protecting Signal Fail (SF) reporting MEP.

APS MEP

Enter the APS PDU handling MEP.

Alarm

Indicates the alarm status on the EPS. A red dot indicates active alarm and green indicates no active alarms.

Buttons

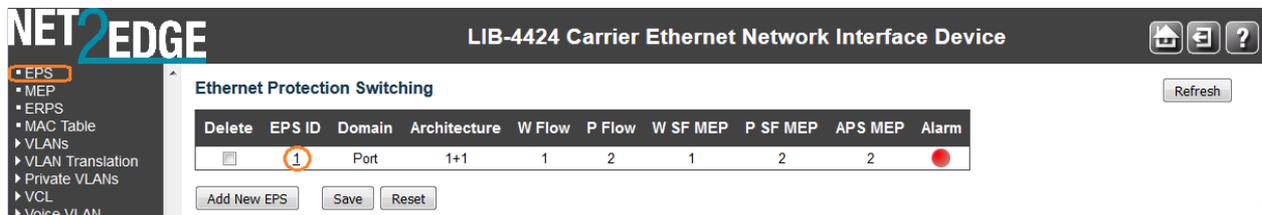
Add New EPS: Click to add a new EPS entry. Note that only one EPS can be added for each save operation.

Save: Click to save changes.

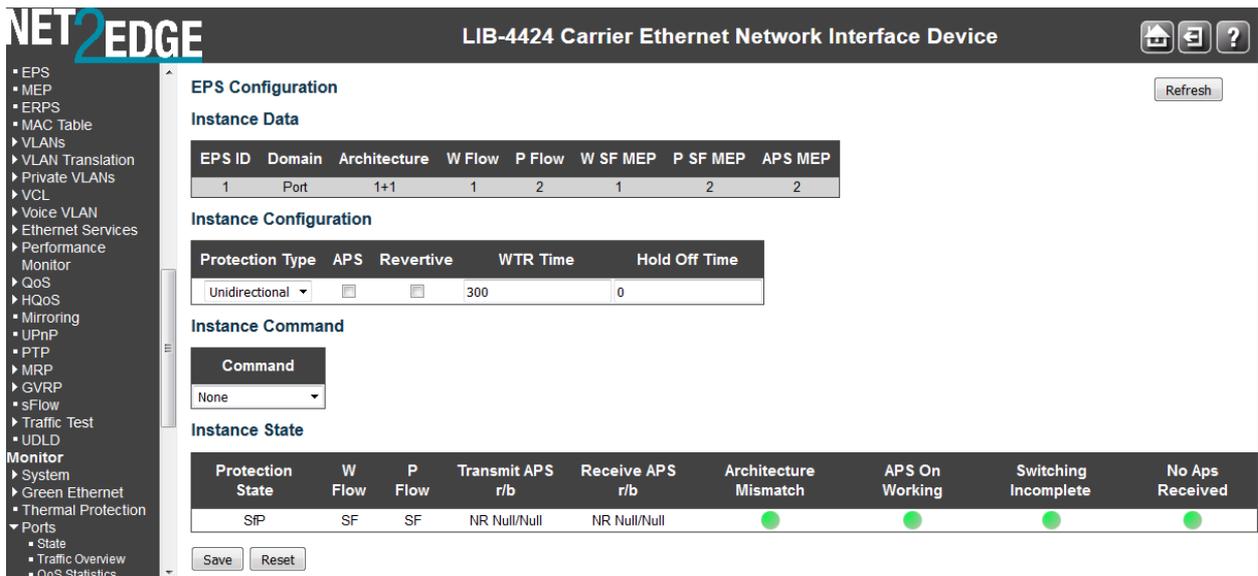
Reset: Click to undo any changes made locally and revert to previously saved values.

EPS Configuration

When you click on the ID of an EPS, its configuration page displays.



For example, if you click on EPS ID 1 on the screen above, the EPS ID 1 configuration screen displays:



This screen lets you configure the EPS Instance Data, Instance Configuration, and Instance Command parameters as well as display the current Instance State.

These parameters are explained below:

EPS Instance Data

This table displays this EPS configuration instance information (EPS ID, Domain, Port, Architecture, W Flow, P Flow, W SF MEP, P SF MEP and APS MEP state). See the descriptions above.

EPS Instance Configuration

If configured, the possible Protection Types are:

- 1+1 Unidirectional, no APS communication.
- 1+1 Unidirectional with APS communication.
- 1+1 Bidirectional with APS communication.
- 1:1 Bidirectional with APS communication.

Configured

Displays green (●) for Up, or red (●) for Down.

Red: This EPS is only created and has not yet been configured (not active).

Green: This EPS is configured (active).

After initial creation of EPS, the EPS instance is deactivated by default, as indicated by the **red** Configured icon in EPS edit mode. There is nothing that immediately indicates how to activate it. You must click on the EPS instance to edit it, and then click **Save** to activate it. When activated, the Configured icon turns **green**.

Protection Type

Select Unidirectional or Bidirectional protection mode:

Unidirectional: EPS in the two ends can select traffic from different working/protecting flow. This is only possible in case of 1+1 protection.

Bidirectional: EPS in the two ends is selecting traffic from the same working/protecting flow. This requires APS enabled. This is mandatory for 1:1 protection.

APS

Check or uncheck the checkbox to enable or disable Automatic Protection Switching (the APS protocol). This is mandatory for 1:1 protection. Check the checkbox to enable the automatic protection switching APS protocol and linear protection switching mechanisms for point-to-point VLAN-based ETH SNC (Sub Network Connection) in Ethernet transport networks per Rec. ITU-T G.8031/Y.1342 (11/2009).

Bidirectional switching always requires APS communication. The only switching type that does not require APS communication is 1+1 unidirectional switching.

Revertive

Check or uncheck the checkbox to enable or disable Revertive mode. The revertive switching to working flow can be enabled or disabled here.

Revertive mode: traffic is restored to the working entities after a switch reason 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 a “Wait to Restore” timer, which is used to avoid chattering of selectors in the case of intermittent defects.

Operationally, in revertive mode, in conditions where working traffic is being received via the protection entity, if local protection switching requests have been previously active and now become inactive, a local WTR state is entered. Since this state now represents the highest priority local request, it is indicated on the transmitted “Request/State” information and maintains the switch. This state normally times out and becomes a NR state after the WTR timer has expired. The WTR timer is deactivated earlier if any local request of higher priority pre-empts this state.

Note that for the decision of whether or not to enter the WTR state, only local requests are considered.

A switch to the protection entity may be maintained by a local WTR state or by a remote request (WTR or other) received via the “Request/State” information. Therefore, in a case where a bidirectional failure for a working entity has occurred and subsequent repair has taken place, the bidirectional reversion back to the working entity does not take place until both WTR timers at both ends have expired.

Non-revertive mode: normal traffic is allowed to remain on the protection entity even after a switch reason has cleared. This is generally accomplished by replacing the previous switch request with a “Do not Revert (DNR)” request, which is low priority.

1+1 protection is often provisioned as non-revertive; the protection is fully dedicated in any case, and this avoids a second “glitch” in the traffic. However there may be reasons to provision this to be revertive (e.g., so that the traffic uses the “short” path except during failure conditions. Certain operator policies also dictate revertive operation even for 1+1).

1:1 protection is usually revertive. It is possible to define the protocol in a way that would permit non-revertive operation for 1:1 protection; however, since the working transport entity is typically more optimized (i.e., in terms of delay and resourcing) than the protection transport entity, it is better to revert and glitch the traffic when the working transport entity is repaired.

In general, the choice of revertive / non-revertive will be the same at both ends of the protection group. However, a mismatch of this parameter does not prevent interworking; it would be peculiar

for one side to go to WTR for clearing of switches initiated from that side, while the other goes to DNR for its switches.

Operationally, in non-revertive mode, in conditions where working traffic is being transmitted via the protection entity, if local protection switching requests have been previously active and now become inactive, a local DNR state is entered. Since this state now represents the highest priority local request, it is indicated on the transmitted “Request/State” information and maintains the switch, thus preventing reversion back to the released bridge-selector position in non revertive mode under NR conditions.

WTR Time

Select either Disabled, 10sec, 30sec, 5min, 6min, 7min, 8min, 9min, 10min, 11min, or 12 minutes as the WTR (Wait To Restore) timing value to be used in revertive switching. For example, set the WTR timer to 5 minutes so you are protected by primary line flapping.

In revertive mode, to prevent frequent operation of the protection switch due to an intermittent defect, a failed working transport entity must become fault-free. After the failed working transport entity meets this criterion, a fixed period of time elapses before a normal traffic signal uses it again. This period, called the wait-to-restore (WTR) period, may be configured in 1 minute steps between 5 and 12 minutes; the default value is 5 minutes.

An SF (or SD, if applicable) condition will override the WTR. To activate the WTR timer appropriately even when both ends concurrently detect clearance of SF, when the local state transits from SF to NR with the requested signal number 1, the previous local state, SF, should be memorized. If both the local state and far-end state are NR with the requested signal number 1, the local state transits to WTR only when the previous local state is SF. Otherwise, the local state transits to NR with the requested signal number 0.

The image shows a dropdown menu titled "WTR Time". The menu is currently open, showing a list of options: Disabled (selected), 10sec, 30sec, 5min, 6min, 7min, 8min, 9min, 10min, 11min, and 12min.

In revertive mode, when the protection is no longer requested (i.e., the failed working transport entity is no longer in SF condition - or SD condition, if applicable, and assuming no other requesting transport entities), a local wait-to-restore state is activated. Since this state becomes the highest in priority, it is indicated on the APS signal (if applicable) and maintains the normal traffic signal from the previously failed working transport entity on the protection transport entity. This state normally times out and becomes a no-request state. The WTR timer deactivates earlier when any higher priority request pre-empts this state.

Hold Off Time

Select Disabled, 100ms, 200ms, 300ms, 400ms, 500ms, 600ms, 700ms, 800ms, 900ms, 1s, 2s, 3s, 4s, 5s, 6s, 7s, 8s, 9s, or 10 seconds as the hold off time to make persistent checks on Signal Fail (SF) before switching. You can set the Hold-off timer to 0 so the switchover to backup happens immediately on connection failure.

A hold-off timer is implemented to coordinate the timing of protection switches at multiple layers or across cascaded protected domains. Its purpose is to allow either a server layer protection switch to have a chance to fix the problem before switching at a client layer, or to allow an upstream protected domain to switch before a downstream domain (e.g., to allow an upstream ring to switch before the downstream ring in a dual node interconnect configuration so that the switch occurs in the same ring as the failure).

The image shows a dropdown menu titled "Hold Off Time". The menu is currently open, showing a list of options: Disabled (selected), 100ms, 200ms, 300ms, 400ms, 500ms, 600ms, 700ms, 800ms, 900ms, 1s, 2s, 3s, 4s, 5s, 6s, 7s, 8s, 9s, and 10s.

Each protection group has a configurable hold-off timer. When a new defect or more severe defect occurs (e.g., a new SF), this event will not be reported immediately to protection switching if the provisioned hold-off timer value is non-zero. Instead, the hold-off timer is started; when it expires, it is checked if a defect still exists on the trail that started the timer. If it does, that defect is reported to protection switching. The defect need not be the same one that started the timer. The hold-off timer applies to both the Working and the Protection transport entities.

Command (EPS Instance Command)

At the **Command** dropdown, select None, Clear, Lock Out, Forced Switch, Manual Switch P, Manual Switch W, Exercise, Freeze, or Lock Out Local.

In general MEF terms, a **Manual Switch** occurs when the network operator switches the network to use the protection resources instead of the working, or vice-versa. By MEF definition, a Manual Switch will not progress to failed resources. Manual switch may occur at any time according to the network operator will, unless the target resource is in failure condition. A **Forced Switch** is when the network operator forces the network to use the protection resources instead of the working resources, or vice-versa, regardless of the state of the resources. A **Lockout** command on a resource makes the resource not available for protection of other resources.

The specific LIB-44xx **Command** dropdown selections are explained below:

None: There is no active local command on this instance. This EPS is only created and has not yet been configured - is not active.

Clear: The active local command will be cleared. This EPS is configured - is active. This clears the active near-end lockout of protection, forced switch, manual switch, WTR state, or exercise command. A Clear is an action, initiated externally, that clears the active external command.

Lock Out: This EPS is locked to working (not active). With 1:N protection (more than one EPS with same protecting flow), when one EPS switches to protecting flow, the other EPS enforces this command

Forced Switch: Forced switch to protecting. This forces normal traffic signal to be selected. A Forced switch-over for normal traffic is a switch-over action, initiated externally, that switches normal traffic to the recovery LSP/span, unless an equal or higher priority switch-over command is in effect.

Manual Switch P: Manual switch to protecting. In the absence of a failure of a working or protection transport entity, forces normal traffic signal to be selected. A Manual switch-over for normal traffic is a switch-over action, initiated externally, that switches normal traffic to the recovery LSP/span, unless a fault condition exists on other LSPs/spans (including the recovery LSP/span) or an equal or higher priority switch-over command is in effect.

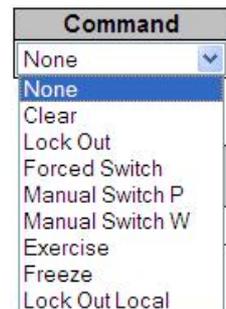
Manual Switch W: Manual switch to working - this is only possible in 1:1 non-revertive.

Exercise: Exercise of the protocol - not traffic effecting. (Exercise of the APS protocol. The signal is chosen so as not to modify the selector.)

Freeze: This EPS is locally frozen, ignoring all input. This freezes the state of the protection group. Until the freeze is cleared, additional near-end commands are rejected. Condition changes and received APS information are ignored. When the freeze command is cleared, the state of the protection group is recomputed based on the condition and received APS information.

A Freeze is a configuration action, initiated externally, that prevents any switch-over action from being taken, and, as such, 'freezes' the current state.

Lock Out Local: This EPS is locally "locked out" - ignoring local SF (signal fail) detected on working.



Instance State (EPS Instance State)

Protection State: displays the current EPS state per the State Transition Tables in the G.8031 standard.

Disabled: Protection state currently disabled.

NoReqW: Currently in No request (NR) on working state (lowest priority).

NoReqP: Currently in No request (NR) on protection state (lowest priority).

Lockout: Currently locked out of protection (LO) state. This EPS is locked to the working (not active) state.

Forced: Currently in Forced switch to protecting state.

SfW: Currently in ‘Signal Fail on Working’ state.

SfP: Currently in ‘Signal Fail on Protection’ state (SF-P; highest priority).

ManualW: Currently in Manual switch to working (MS-W) state.

ManualP: Currently in Manual switch to protection state.

Wtr: Currently in Wait to Restore (WTR) state.

ExerW: Currently in Exercise (EXER) working state.

ExerP: Currently in Exercise (EXER) protection state.

RevReqW: Currently in Reverse request (RR) working state.

RevReqP: Currently in Reverse request (RR) protection state.

DoNotRev: Currently in Do not revert (DNR) state.

Idle: no detected failure, no active automatic or external command, and receiving only "NR, RB" R-APS messages.

W Flow: Displays Working flow status of **OK**, **SF**, or **SD**, where:

OK: State of working flow is ok

SF: State of working flow is Signal Fail

SD: State of working flow is Signal Degrade (for future use)

P Flow: Displays Protection flow status of **OK**, **SF**, or **SD**, where:

OK: State of protecting flow is ok

SF: State of protecting flow is Signal Fail

SD: State of protecting flow is Signal Degrade (for future use)

Transmit APS r/b: Displays the transmitted APS according to the State Transition Tables in the G.8032 standard. In this field, **RB** indicates ‘RPL Blocked’, and **NR** indicates ‘No Request’ (e.g., **NR Null/Null** displayed). No request (NR) is the ring protection condition when no local protection switching requests are active. **EXER Null/Normal** indicates Exercise (EXER) protection state null / normal state.

Receive APS r/b: Displays the received APS according to the State Transition Tables in the G.8032 standard. In this field, **rb** indicates ‘RPL Blocked’, and **NR** indicates ‘No Request’ (e.g., **NR Null/Null** displayed). The status when RPL is blocked (used by RPL Owner in NR).

Architecture Mismatch: Indicates whether the architecture indicated in the received APS does not match the locally configured architecture. Displays a green dot for Up, or a red dot for Down. With all of the options for provisioning, there are opportunities for mismatches between the provisioning at the two ends. These provisioning mismatches take one of several forms, such as Mismatches where proper operation is not possible, or Mismatches where one or both sides can adapt their operation to provide a degree of interworking in spite of the mismatch, or Mismatches that do not prevent interworking. An example is the revertive / non-revertive mismatch.

APS on working: Indicates whether the APS is received on the working flow. Displays a green dot for Up, or a red dot for Down.

Instance State

Protection State
SfP

Switching Incomplete: Indicates whether the Traffic is not selected from the same flow instance in the two ends. Displays a green dot for Up, or a red dot for Down.

No APS Received: Indicates if the APS PDU is received from the other end. Displays a green dot for Up, or a red dot for Down.

Note: Since the protection switching mechanism requires monitoring for both working and protection transport entities, MEPs must be activated for monitoring the working and protection transport entities. See the SOAM “[MEP Configuration](#)” section below for more information.

MEP Configuration

LIB-44xx MEP (Maintenance Entity Group End Point) configuration is done from the **Configuration > MEP** menu path.

A MEP (MEG End Point) is an endpoint in a Maintenance Entity Group (per ITU-T Y.1731).

Service OAM Standards

The term “Service OAM” is commonly used for the ITU-T Y.1731, IEEE802.1ag, and MEF 17 standards, all covering OAM (Operation, Administration and Maintenance). These standards were developed to cover monitoring and error detection, which are missing in standard Ethernet.

Monitoring and error detection are key weaknesses in Ethernet compared to SONET and SDH, both of which have monitoring built in on multiple levels, with each level having parity checking and trace identifiers. Ethernet has CRC checking on each transmitted frame, but if no frames are transmitted, there is no way to quickly detect a link failure, and no way to detect failures on a path. This is the reason Service OAM standards were developed.

Service OAM contains a suite of OAM functions which can be divided into two main groups: Fault Management and Performance Management.

All of the OAM functions are listed below:

OAM Functions for Fault Management

Ethernet Continuity Check (ETH-CC)

Ethernet Loopback (ETH-LB)

Ethernet Link Trace (ETH-LT)

Ethernet Alarm Indication Signal (ETH-AIS)

Ethernet Remote Detect Indication (ETH-RDI)

Ethernet Locked Signal (ETH-LCK)

Ethernet Test Signal (ETH-Test)

Ethernet Automatic Protection Signal (ETH-APS)

Ethernet Maintenance Communication Channel (ETH-MCC)

Ethernet Experimental OAM (ETH-EXP) (Y.1731 only)

Ethernet Vendor Specific OAM (ETH-VSP)

OAM Functions for Performance Management (Y.1731 only)

Frame Loss Measurement (ETH-LM)

Frame Delay Measurement (ETH-DM)

Throughput Measurement

Y.1731 is prepared in cooperation with IEEE802.1ag, shares the same basic PDU (Packet Data Unit) format and the fault management is the same, but Y.1731 also contains performance monitoring

which is not covered by the current IEEE standard. Since Y.1731 is the most comprehensive standard, the content and terms used in this document are based on Y.1731.

Basic OAM Functions and Terms.

The OAM functions are based on transmission and reception of OAM frames (i.e., PDU frames). OAM frames are exchanged within a Maintenance Entity (ME) and the points that transmit and receive OAM frames are called Maintenance Entity Group End Points (MEPs).

The ME Group has a unique ID and each MEP has a unique ID within the MEG. OAM frames have a unique EtherType of 0x8902 and are transmitted either as Unicast or Multicast within a dedicated range of MAC addresses (01-80-C2-00-00-30 and 01-80-C2-00-00-3F).

ME groups can be nested but cannot overlap. To accommodate nesting, the OAM frame contains a MEG level (i.e., a MEP at a certain level will forward OAM frames of a higher level and block OAM frames at a lower level). The MEG levels are divided into three roles: the 'Customer' role is assigned three MEG levels (7, 6, and 5), the 'Provider' role is assigned two MEG levels (4 and 3), and the 'Operator' role is assigned three MEG levels (2, 1, and 0).

Especially with Fault Protection, it is possible to have a MEG Intermediate Point (MIP). A MIP reacts only to link trace and Unicast Loopback PDUs, and forwards all OAM frames.

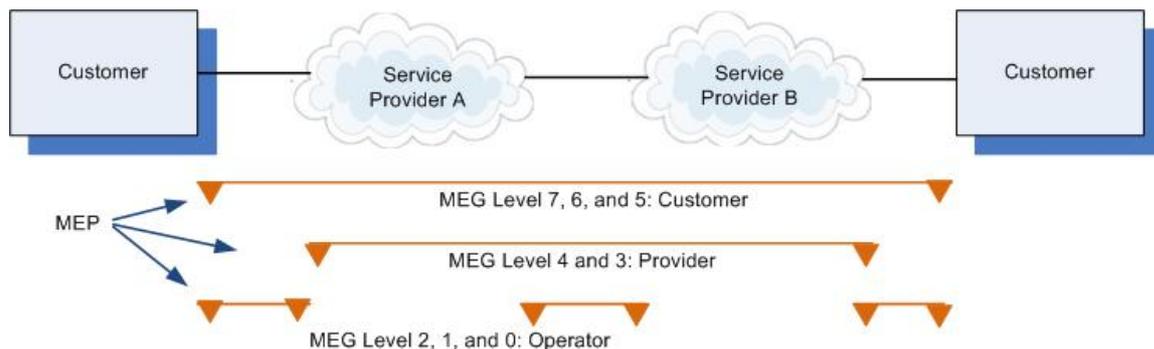


Figure 1. MEG Levels in the Network

The LIB-44xx implements Service OAM (SOAM), compliant with IEEE802.1ag and ITU Y.1731 standards. SOAM operates over the LIB-44xx VLAN configuration. The LIB-44xx SOAM features include:

All 802.1ag and Y.1731 managed objects such as MD, MA, MEG, MEP and MIP can be created, modified and deleted from the LIB-44xx through all management interfaces. All eight MEG levels are supported on a VLAN (C-Tagged, S-tagged and double Tagged) and per-port basis. Both UP and Down MEPs are supported. Both 802.1ag and Y.1731 type of MIPs are supported.

The SOAM continuity check is supported with interval values of 100 ms, 3.33 ms, 1 second, 10 seconds, 1 minute and 10 minutes. Statistics for CCM sent/received are available. The CCM database is maintained for each MEP. The initial LIB-44xx release supports only point-to-point deployment of SOAM sessions. All CCM errors such as remoteCCM, RDI, MACStatus, errorCCM, and crossConnect are reported in MEP status, and SNMP traps are raised for errors.

The Link Trace protocol is supported and the destination MP can be specified using MAC address or MEP IDs. Link trace replies (LTRs) from all MPs in the path to reach the final MP are available on the management interface.

The Loopback protocol is supported and the destination MEP can be specified using MAC address or MEP IDs. The loopback can be a Unicast or a Multicast message.

The Y.1731 AIS and LCK protocol is supported for fault monitoring and isolation. SNMP traps are raised for the faults.

Performance monitoring for frame delay and frame loss are supported. Delay measurement is available on-demand per the Y.1731 standard.

The LIB-44xx supports a maximum of 10 MEPS per port assignable to any VLAN(s) on a port and a minimum of 80 MEPS per LIB-44xx. The maximum number of MEPs at each CCM interval that can be supported without affecting performance depends on what else is configured on the LIB-44xx. As per the MEF UNI requirements, each port needs to support a minimum of 8 VLANs and as per SOAM standards each VLAN can have a Max of 8 MEPs (one at each level). The LIB-44xx system limitation is 80 MEPs, but on a per-port basis it supports 16 VLANs with 8 MEPS per VLAN, but not on all ports at the same time.

Ethernet Service (OAM) standards are designed to simplify the management of Carrier Ethernet services with end-to-end service visibility, fault isolation, reporting and continuous performance monitoring as specified in IEEE 802.1ag, ITU-T Y.1731, and MEF (Metro Ethernet Forum) specifications and standards.

The differences between related terms used by IEEE 802.1ag, ITU-T Y.1731, and MEF are shown below:

Difference in IEEE 802.1ag, ITU-T 1731, and MEF Terms

The SOAM MEP terms vary slightly between the IEEE 802.1ag, ITU-T 1731, and MEF standards. The differences are noted below:

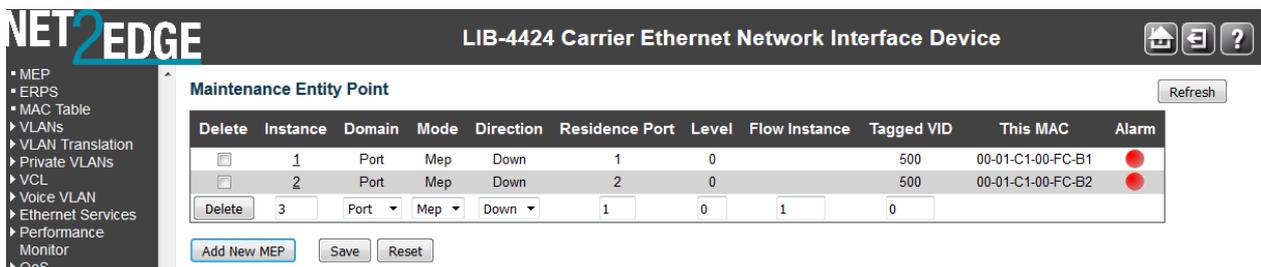
Table 1: SOAM Terms - IEEE 802.1ag vs. ITU-T 1731 vs. MEF

IEEE 802.1ag Term	ITU-T Y.1731 Term	MEF Term*
Maintenance Entity Group (MEG)	Maintenance Association (MA)	Maintenance Entity Group (MEG)
Maintenance Entity Group Identifier (MEG ID)	Maintenance Association Identifier (MAID)	Both MEG and MAID
No IEEE equivalent of this ITU term.	Maintenance Domain (MD)	No MEF equivalent of this ITU term. MEF 30 uses MD only in describing the format of a MAID.
Maintenance Entity Group Level (MEG level)	Maintenance Domain Level (MD Level)	MEF 30 uses MEG level
ME (Maintenance Entity)	ME (Maintenance Entity)	ME (Maintenance Entity) per MEF 17, 30, 31.
MA (Maintenance Association)	MEG (ME Group)	MA Maintenance Association (equivalent to a 'MEG'). See MEF 30.
MAID (MA Identifier)	MEGID (MEG Identifier)	MAID (Maintenance Association

		Identifier) (equivalent to a MEG ID). See MEF 30.
MD (Maintenance Domain)	(No such construct available)	MD Maintenance Domain (equivalent to "OAM Domain" in MEF 17). Also see MEF 30, 31.
MD Level	MEG Level	MD Level Maintenance Domain Level (equivalent to "MEG level"). See also MEF 30.
MEP (MA End Point)	MEP (MEG End Point)	MEP (MEG End Point) per MEF 17. See also MEF 30, 31.
MIP (MD Intermediate Point)	MIP (MEG Intermediate Point)	MIP (MEG Intermediate Point) per MEF 17. See also MEF 30, 31.
(No such construct available)	Server MEP	No MEF equivalent of this term.

* See the MEF Glossary for a current summary of the terms used in MEF specifications at http://metroethernetforum.org/page_loader.php?p_id=147. Refer to each individual context and the individual glossaries in each MEF specification. See the MEF Technical Specifications at http://metroethernetforum.org/page_loader.php?p_id=29.

From the **Configuration > MEP** menu path, click the **"Add New MEP"** button to display the table. The Maintenance Entity Point (MEP) instances are configured here.



The MEP table parameters are explained below:

Delete

This checkbox is used to mark a MEP for deletion in the next Save operation.

Instance

The ID of the MEP. Click on the ID of a MEP to enter the configuration page (see description below).

Domain

Port: This is a MEP in the Port Domain. 'Flow Instance' is a Port.

Evc: This is a MEP in the EVC Domain. 'Flow Instance' is an EVC. This selection is required for MIP configuration.

Mode

Mep: This is a Maintenance Entity End Point.

Mip: This is a Maintenance Entity Intermediate Point. This **MIP** configuration requires **EVC** to be selected in the **Domain** field (above).

Direction

The LIB-44xx MEP direction naming conventions include Down/Ingress and Up/Egress.

Down/Ingress: This is an Ingress/Down MEP - monitoring ingress traffic on the 'Residence Port'.

Up/Egress: This is an Egress/Up MEP - monitoring egress traffic on the 'Residence Port'.

Residence Port

The port which a MEP is monitoring - see the 'Direction' parameter description above.

Level

The MEG level of this MEP (0-7). The defaults per [MEF 30](#) are:

MEG	Default MEG Level	Suggested Use (MEF 30)
Subscriber MEG	6	Subscriber monitoring of an Ethernet service.
Test MEG	5	SP isolation of subscriber reported problems.
EVC MEG	4	SP monitoring of provided service.
Service Provider MEG	3	SP Monitoring of Service Provider network.
Operator MEG	2	Netw. Operator monitoring of the portion of a network.
UNI MEG	1	Service Provider monitoring of a UNI.
ENNI MEG	1	Network Operators' monitoring of an ENNI.

(where SP = Service Provider)

Note: Assignment of numerical MEG Levels to 'subscriber' (or customer) role, Service Provider role, and Operator role is somewhat arbitrary since those terms imply business relationships that cannot be standardized. For example, a 'subscriber' (or customer) may also be an Operator seeking a service from another Operator. The MEG Level default values are consistent with a shared MEG Level model across Subscriber, Operators, and Service Providers.

Note: The MEF and Broadband Forum (BBF) are not aligned on the use of MEG Level 5. If interworking between an MEF compliant implementation and a BBF compliant implementation is required, an agreement on the use of MEG Level 5 is required between the two parties.

Flow Instance

The MEP is related to this flow - See 'Domain' above.

Tagged VID

Port MEP: An outer C/S-tag (depending on VLAN Port Type) is added with this VID (e.g., **4**).

A '0' displayed means no Tag to be added with this VID.

This MAC

The MAC address (e.g., 00-C0-F2-00-00-02) of this MEP, which can be used by another MEP when unicast is selected (info only).

Alarm

There is an active alarm on the MEP. Red LED ● = Down, Green LED ● = Up.

Note: Click the **Refresh** button to verify the status at the end of the configuration process.

Buttons

Add new MEP: Click to add a new MEP entry.

Refresh: Click to refresh the page immediately.

Save: Click to save changes. Only one MEP can be added per Save operation.

Reset: Click to undo any changes made locally and revert to previously saved values.

An LIB-44xx SOAM configuration with five MEP instances defined is shown below:

Delete	Instance	Domain	Mode	Direction	Residence Port	Level	Flow Instance	Tagged VID	This MAC	Alarm
<input type="checkbox"/>	1	Port	Mep	Down	1	0		500	00-01-C1-00-FC-B1	●
<input type="checkbox"/>	2	Port	Mep	Down	2	0		500	00-01-C1-00-FC-B2	●
<input type="checkbox"/>	3	VLAN	Mep	Up	3	0	300		00-01-C1-00-FC-B3	●
<input type="checkbox"/>	4	Port	Mep	Down	4	0		1000	00-01-C1-00-FC-B4	●

You can view and configure each MEP instance as discussed below:

MEP Instance Configuration

When you click on the Instance of an MEP in the table, its configuration page displays.

For example, if you click on **1** in the Instance column on the screen above, the MEP Configuration screen displays for MEP Instance 1:

MEP Configuration

Instance Data

Instance	Domain	Mode	Direction	Residence Port	Flow Instance	Tagged VID	EPS Instance	This MAC	Oper State
1	Port	Mep	Down	1		500	1	00-01-C1-00-FC-B1	Up

Instance Configuration

Level	Format	Domain Name	MEG id	MEP id	Tagged VID	VOE	Syslog	cLevel	cMEG	cMEP	cAIS	cLCK	cLoop	cConfig	cDEG	cSSF	aBLK	aTSD	aTSF
0	ITU ICC		ICC000MEG0000	1	500	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>											

Peer MEP Configuration

Delete	Peer MEP ID	Unicast Peer MAC	cLOC	cRDI	cPeriod	cPriority
No Peer MEP Added						

Functional Configuration

Continuity Check

Enable	Priority	Frame rate	TLV	Enable	Priority	Cast	Type	Last Octet
<input type="checkbox"/>	0	1 f/sec	<input type="checkbox"/>	<input type="checkbox"/>	0	Multi	L-APS	1

TLV Configuration

Organization Specific TLV (Global)

OUI First	OUI Second	OUI Third	Sub-Type	Value
0	0	12	1	2

TLV Status

Peer MEP ID	CC Organization Specific					CC Port Status		CC Interface Status		
	OUI First	OUI Second	OUI Third	Sub-Type	Value	Last RX	Value	Last RX	Value	Last RX

Link State Tracking

Enable

This screen lets you configure the MEP Instance Data, Instance Configuration, and Functional Configuration parameters. Here you can add a new peer MEP and configure Fault Management and/or Performance Monitoring.

These parameters are explained below:

MEP Instance Data

Displays this MEP configuration instance information (MEP Instance, Domain, Mode, Direction, Residence Port, Flow Instance, Tagged VID, EPS Instance, and This MAC). See the descriptions above.

MEP Instance Configuration

EVC Policy ID

This is relevant for an EVC Egress/Up-MEP. This is the Policy number of the relevant ECE (0-255). The Policy ID that the generated TST frames will get as IS1 action. It can be the same as any ECE Policy number, enabling it to hit the same ACL and thereby the same EVC policer.

Level

Select the MEG level for this MEP (0-7) from the dropdown.

Format

This is the configuration of the two possible Maintenance Association Identifier formats.

ITU ICC: This is defined by ITU. The 'ICC' can be up to 6 characters long. The 'MEG ID' can be up to 7 characters long.

IEEE String: This is defined by IEEE. The 'Domain Name' can be up to 8 characters long. The 'MEG ID' can be up to 8 characters long.

ICC/Domain Name

This is either ITU ICC (MEG ID value = 1-6 characters) or IEEE Maintenance Domain Name - depending on 'Format'. See the 'Format' description above. The default is 'TRNSTN'.

MEG Id

Displays either an ITU UMC (MEG ID value 7-13) or an IEEE Short MA Name - depending on 'Format'. See the 'Format' description above. In case of ITU ICC format this can be up to 7 characters.

If only 6 characters are entered, the MEG ID value 13 will become NULL.

This section uses the Maintenance Association (MA) and Maintenance Association Identifier (MAID) terminology of [IEEE 802.1ag] to clarify the formatting of the MEG ID / MAID. As specified per [IEEE 802.1ag], a MAID has two components consisting of the MD Name and the Short MA Name.

Per MEF 30, the MEG ID must be unique within a MEN, operator's network, where an operator and customer connect, or where two operators interconnect. When a MEG has MEPs or MIPs in more than one network (which is true for all MEGs other than the Operator MEG, and other than a Subscriber MEG with no MIPs configured), then all involved parties must agree to the naming format.

For an ENNI MEG, the MEG ID / MAID must have a format and a value that are jointly agreed upon by the providers of both ends of the ENNI.

MEP Id

This value will become the transmitted two-byte CCM MEP ID.

Tagged VID

Enter a VLAN ID from **0-4094**. Entering '0' means no Tag added.

cLevel

Fault Cause indicating that a CCM is received with a lower level than the configured for this MEP. Green dot = up, Red dot = Down.

cMEG

Fault Cause indicating that a CCM is received with a MEG ID different from configured for this MEP. Green dot = up, Red dot = Down.

cMEP

Fault Cause indicating that a CCM is received with a MEP ID different from all 'Peer MEP ID' configured for this MEP. Green dot = up, Red dot = Down.

cAIS

Fault Cause indicating that an AIS PDU is received (Alarm Indication Signal). Green dot = up, Red dot = Down.

cLCK

Fault Cause indicating that a LCK PDU is received. Green dot = up, Red dot = Down.

cSSF

Fault Cause indicating that the server layer is indicating Signal Fail (SF). Green dot = up, Red dot = Down.

aBLK

The consequent action of blocking service frames in this flow is active. Green dot = up, Red dot = Down.

aTSF

The consequent action of indicating Trail Signal Fail towards protection is active. Green dot = up, Red dot = Down.

Delete

This checkbox is used to mark a Peer MEP for deletion in next Save operation.

Peer MEP ID

This value will become an expected MEP ID in a received CCM - see 'cMEP' above.

Unicast Peer MAC

This MAC will be used when unicast is selected with this peer MEP. Also, this MAC is used to create hardware checking of receiving CCM PDU (LOC detection) from this MEP.

cLOC

Fault Cause indicating that no CCM has been received (in 3,5 periods) from this peer MEP.

cRDI

Fault Cause indicating that a CCM is received with Remote Defect Indication from this peer MEP.

cPeriod

Fault Cause indicating that a CCM is received with a period different what is configured for this MEP - from this peer MEP.

cPriority

Fault Cause indicating that a CCM is received with a priority different what is configured for this MEP - from this peer MEP.

MEP Functional Configuration**Continuity Check**

Continuity Check based on transmitting/receiving CCM PDU can be enabled/disabled.

Continuity Check Messages (CCMs) are 'heartbeat' messages exchanged periodically between the MEPs in a service. Each MEP sends out multicast CCMs, and receives CCMs from all the other MEPs in the service (its 'peer' MEPs). This allows each MEP to discover its peer MEPs and to verify that there is connectivity between them. MIPs also receive CCMs; the MIPs use the discovered information to build a MAC learning database for use when responding to a Linktrace.

Enable

Continuity Check based on transmitting/receiving CCM PDU can be enabled/disabled. The CCM PDU is always transmitted as Multi-cast Class 1.

Priority

The priority to be inserted as PCP bits in TAG (if any). In case of enable of Continuity Check and Loss Measurement both implemented on SW based CCM, 'Priority' must be the same.

Frame rate

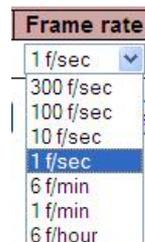
Select the frame rate of CCM PDU. The selections are 1 f/sec, 300 f/sec, 100 f/sec, 10 f/sec, 1 f/sec, 6 f/min, 1 f/min, or 6 f/hour. This is the inverse of transmission period as described in Y.1731. This value has the following uses:

The transmission rate of the CCM PDU.

Fault Cause cLOC is declared if no CCM PDU has been received within 3.5 periods - see 'cLOC'.

Fault Cause cPeriod is declared if a CCM PDU has been received with a different period - see 'cPeriod'.

Selecting 300f/sec or 100f/sec will configure HW based CCM (if possible). Selecting other frame rates will configure SW based CCM. In case of enable of Continuity Check and Loss Measurement both implemented on SW based CCM, 'Frame Rate' must be the same.

**APS Protocol**

APS information is carried within the APS OAM PDU, which is one of a several Ethernet OAM PDUs. OAM PDU formats for each type of Ethernet OAM operation are as defined in Y.1731.

Enable

Automatic Protection Switching protocol information transportation based on transmitting / receiving R-APS/L-APS PDU can be enabled or disabled. This must be enabled to support ERPS/ELPS implementing APS. This is only valid with one Peer MEP configured.

Priority

The priority to be inserted as PCP bits in TAG (if any).

Cast

Selection of APS PDU transmitted unicast or multi-cast. The unicast MAC will be taken from the 'Unicast Peer MAC' configuration. Unicast is only valid for L-APS - see 'Type' below: The R-APS PDU is always transmitted with multi-cast MAC described in G.8032.

Type

R-APS: APS PDU is transmitted as R-APS; this is for ERPS. (Automatic Protection Switching Protocol Data Unit transmitted as Ring APS protocol per Rec. ITU-T G.8032/Y.1344 (03/2010) for Ethernet Ring Protection Switching.)

L-APS: APS PDU is transmitted as L-APS; this is for ELPS (G.8031 Ethernet Linear Protection Switching).

Last Octet

This is the last octet of the transmitted and expected RAPS multi-cast MAC. In G.8031 (03/2010) a RAPS multi-cast MAC is defined as 01-19-A7-00-00-XX. In the current standard, the value for this last octet is '01' and the usage of other values is for further study.

Buttons

Fault Management: Click to go to the Fault Management page (see below).

Performance Monitor: Click to go to the Performance Monitor page (see below).

Refresh: Click to refresh the page immediately.

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Fault Management (FM)

This page lets you view and configure the Fault Management of the current MEP Instance. Note that after a system reboot, the MEP PM and FM become disabled.

NET2EDGE LIB-4424 Carrier Ethernet Network Interface Device

Fault Management - Instance 1 - MEP id 1 Refresh

Loop Back

Enable	DEI	Priority	Cast	Peer MEP	Unicast MAC	To Send	Size	Interval
<input checked="" type="checkbox"/>	<input type="checkbox"/>	0	Multi	1	00-00-00-00-00-00	10	64	100

Loop Back State

Transaction ID	Transmitted	Reply MAC	Received	Out Of Order
No Replies				

Link Trace

Enable	Priority	Peer MEP	Unicast MAC	Time To Live
<input type="checkbox"/>	0	1	00-00-00-00-00-00	1

Link Trace State

Transaction ID	Time To Live	Mode	Direction	Forwarded	Relay	Last MAC	Next MAC
No Transactions							

Test Signal

Tx	Rx	DEI	Priority	Peer MEP	Rate	Size	Pattern	Sequence Number
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	0	1	1000	64	All Zero	<input type="checkbox"/>

Test Signal State

TX frame count	RX frame count	RX rate	Test time	Clear
0	0	0	0	<input type="checkbox"/>

Client Configuration

	Flow									
Domain	VLAN									
Instance	0	0	0	0	0	0	0	0	0	0
Level	0	0	0	0	0	0	0	0	0	0
AIS prio	0	0	0	0	0	0	0	0	0	0
LCK prio	0	0	0	0	0	0	0	0	0	0

AIS

Enable	Frame Rate	Protection
<input type="checkbox"/>	1 f/sec	<input type="checkbox"/>

LOCK

Enable	Frame Rate
<input type="checkbox"/>	1 f/sec

Back Save Reset

Loop Back

Enable

Loop Back based on transmitting/receiving LBM/LBR PDU can be enabled or disabled. Loop Back is automatically disabled when all 'To Send' LBM PDUs have been transmitted - waiting 5 seconds for all LBR from the end.

Loopback Messages (LBMs) and Loopback Replies (LBRs) are used to verify connectivity between a local MEP and a particular remote MP. At the request of the administrator, a local MEP sends unicast LBMs to the remote MP. On receiving each LBM, the target maintenance point sends an LBR back to the originating MEP. Loopback indicates whether the destination is reachable or not; it does not allow hop-by-hop discovery of the path; it is conceptually similar to an ICMP Echo (ping).

Dei

The DEI to be inserted as PCP bits in a Tag (if any). The DEI (Drop Eligible Indicator) is a 1-bit field in the VLAN tag.

Priority

The priority to be inserted as PCP bits in a Tag (if any).

Cast

Selection of LBM PDU transmitted unicast or multi-cast. The unicast MAC will be configured through 'Peer MEP' or 'Unicast Peer MAC'. Towards MIP only unicast Loop Back is possible. Select 'Uni' or 'Multi'.

Peer MEP

This is only used if the 'Unicast MAC' is configured to all zeros. The LBM unicast MAC will be taken from the 'Unicast Peer MAC' configuration of this peer.

Unicast MAC

This is only used if NOT configured to all zeros. This will be used as the LBM PDU unicast MAC. This is the only way to configure Loop Back towards a MIP.

To Send

The number of LBM PDU to send in one loop test. The valid range is 1-1000.

Size

The number of bytes in the LBM PDU Data Pattern TLV. The valid range is 1-1400.

Gap

The gap between transmitting LBM PDU in 10ms increments. The valid range is 0-100, where '0' is as fast as possible).

Loop Back State

Transaction ID

The transaction id of the first LBM transmitted. For each LBM transmitted (To Send) the transaction id in the PDU is incremented.

Reply MAC

The MAC of the replying MEP/MIP. In case of multi-cast replies can be received from all peer MEP in the group.

Received

The total number of LBR PDU received from this 'Reply MAC'.

Out Of Order

The number of LBR PDU received from this 'Reply MAC' with incorrect 'Transaction ID'.

Link Trace

Enable

Link Trace based on transmitting/receiving LTM/LTR PDU can be enabled/disabled. Link Trace is automatically disabled when all 5 transactions are done with 5 sec. interval - waiting 5 sec. for all LTR in the end. The LTM PDU is always transmitted as Multi-cast Class 2.

Linktrace Messages (LTMs) and Linktrace Replies (LTRs) are used to track the path, hop-by-hop, to a unicast destination MAC address. At the request of the operator, a local MEP sends an LTM. Each

hop that has a maintenance point sends an LTR back to the originating MEP. This allows the administrator to discover connectivity data about the path (conceptually similar to IP traceroute).

Priority

The priority to be inserted as PCP bits in TAG (if any).

Peer MEP

This is only used if the 'Unicast MAC' is configured to all zeros. The Link Trace Target MAC will be taken from the 'Unicast Peer MAC' configuration of this peer.

Unicast MAC

This is only used if NOT configured to all zeros. This will be used as the Link Trace Target MAC. This is the only way to configure a MIP as Target MAC.

Time To Live

This is the LTM PDU TTL value as described in Y.1731. This value is decremented each time forwarded by a MIP. Will not be forwarded on reaching zero. The valid LT TTL range is 0-255.
Link Trace State

Transaction ID

The transaction ID is incremented for each LTM send. This value is inserted the transmitted LTM PDU and is expected to be received in the LTR PDU. Received LTR with wrong transaction id is ignored. There are five transactions in one Link Trace activated.

Time To Live

This is the TTL value taken from the LTM received by the MIP/MEP sending this LTR - decremented as if forwarded. The valid LT TTL range is 0-255.

Mode

Indicates if it was a MEP/MIP sending this LTR.

Direction

Indicates if MEP/MIP sending this LTR is ingress/egress.

Relayed

Indicates if MEP/MIP sending this LTR has relayed/forwarded the LTM.

Last MAC

The MAC identifying the last sender of the LBM causing this LTR - initiating MEP or previous MIP forwarding.

Next Mac

The MAC identifying the next sender of the LBM causing this LTR - MIP forwarding or terminating MEP.

Test Signal

The Ethernet test signal function (ETH-Test) is used to perform one-way on-demand in-service or out-of-service diagnostics tests. This includes verifying bandwidth throughput, frame loss, bit errors, etc.

When configured to perform such tests, a MEP inserts frames with ETH-Test information with specified throughput, frame size and transmission patterns. When an out-of-service ETH-Test function is performed, client data traffic is disrupted in the diagnosed entity.

When configured to perform such tests, a MEP inserts frames with ETH-Test information with specified throughput, frame size, and transmission patterns. A test signal generator associated with a MEP can transmit TST frames as often as the test signal generator configuration. When a MEP receives TST frames, it examines them to ensure that the MEG Level corresponds to its own configured MEG Level. If the receiving MEP is configured for the ETH-TST function, the test signal detector associated with the MEP detects bit errors from the pseudo-random bit sequence of the received TST frames and reports such errors.

Enable

Test Signal based on transmitting TST PDU can be enabled or disabled here.

Dei

The DEI to be inserted as PCP bits in Tag (if any).

Priority

The priority to be inserted as PCP bits in Tag (if any).

Peer

This is only used if the 'Unicast MAC' is configured to all zero. The TST frame destination MAC will be taken from the 'Unicast Peer MAC' configuration of this peer.

Unicast MAC

This is only used if not configured to all zero. This will be used as the TST frame destination MAC.

Rate

The TST frame transmission bit rate in megabits per second (Mbps). The valid range is 1-400.

MEP Size

The TST frame size. This is entered as the wanted size (in bytes) of a un-tagged frame containing TST OAM PDU - including CRC (four bytes).

For example, when 'Size' = 64 => Un-tagged frame size = DMAC(6) + SMAC(6) + TYPE(2) + TST PDU LENGTH(46) + CRC(4) = 64 bytes

The transmitted frame will be four bytes longer for each tag added - 8 bytes in case of a tunnel EVC.

The transmitting frame rate will be adjusted according to the actually transmitted frame size to obtain correct transmission bit rate. The valid range is 1-1518.

Pattern

Select **All Zero**, **All One**, or **10101010**. The 'empty' TST PDU has the size of 12 bytes. In order to achieve the configured frame size a data TLV will be added with a pattern.

For example when 'Size' = 64 => Un-tagged frame size = DMAC(6) + SMAC(6) + TYPE(2) + TST PDU LENGTH(46) + CRC(4) = 64 bytes.

The TST PDU must be 46 bytes, so a pattern of 46-12=34 bytes will be added.

All Zero: Pattern will be '00000000' (all 0s).

All One: Pattern will be '11111111' (all 1s).

10101010: Pattern will be '10101010' (alternating 1s and 0s).

Sequence Number

Check or uncheck to enable or disable ETH-TST sequence numbering. Each test frame has a sequence number. A MEP cannot repeat sequence numbers within one minute. Be sure the bandwidth used by the test frames is limited.

TX frame count

The number of transmitted TST frames since the last 'Clear'.

RX frame count

The number of received TST frames since the last 'Clear'.

RX rate

The current received TST frame bit rate in **100** Kbps. This is calculated on a 1 second basis, starting when first TST frame is received after 'Clear'. The frame size used for this calculation is the first received after 'Clear'.

Test time

The number of seconds passed since first TST frame received after last 'Clear'.

Clear

This will clear all Test Signal State. Transmission of TST frame will be restarted. Calculation of 'Rx frame count', 'RX rate' and 'Test time' will be started when receiving first TST frame.

Client Configuration

Domain

The domain of the client layer. Must be **EVC**. Only a **Port** MEP is able to be a server MEP having relation to a client layer.

Level

Client layer level - meaning that PDU transmitted in the client layer flows will be on this level.

Flow

Client layer flow instance numbers; enter a maximum of 10 instances. The valid Flow range is 0-65535.

AIS (Alarm Indication Signal)

Alarm Indication Signal (AIS) messages are used to rapidly notify MEPs when a fault is detected in the middle of a domain, in an event driven way. With AIS, MEPs can learn of a fault much sooner than if they rely on detecting a loss of continuity, etc.

Enable

Insertion of AIS (AIS PDU transmission) in client layer flows can be enabled or disabled here.

Prio

The priority to be inserted as PCP bits in TAG (if any).

Frame Rate

Select the frame rate of AIS PDU. This is the inverse of transmission period as described in Y.1731. Select **1f/second** or **1f/minute**.

Protection

Checking this checkbox causes the first three AIS PDUs to be transmitted as fast as possible - in case of using this for protection in the end point.

LOCK

The Ethernet Locked Signal (ETH-LCK) is used to block reaction to a fault situation (much like the ETH-AIS is used to distribute fault conditions). ETH-LCK is normally used in test situations where a change to the network should not result in a protection switch, for example.

Enable

Insertion of LCK (LCK PDU transmission) in client layer flows can be enable or disabled here.

Priority

The priority to be inserted as PCP bits in TAG (if any).

Frame Rate

Select the frame rate of LCK PDU. This is the inverse of transmission period as described in Y.1731. Select **1f/second** or **1f/minute**.

Buttons

Back: Return to the “MEP Configuration” page.

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Refresh: Click to refresh the page immediately.

MEP Performance Monitoring (PM)

This page lets you view and configure the current MEP Instance performance monitoring in terms of LM (Loss Measurement) and DM (Delay measurement). Note that after a system reboot, the MEP PM and FM become disabled.

Service provider SLAs depend on the ability to measure and monitor performance metrics for packet loss and one-way and two-way delay, plus related metrics such as delay variation. This measurement ability also provides operators with better visibility into network performance characteristics, thus facilitating planning, troubleshooting, and overall network performance evaluation.

Performance Monitor - Instance 1 - MEP id 1

Performance Monitoring Data Set

Enable

Loss Measurement

Enable	Priority	Frame rate	Cast	Ended	FLR Interval	Flow Counting	Oam Counting	Loss Threshold
<input type="checkbox"/>	0	1 f/sec	Multi	Single	5	<input type="checkbox"/>	Y1731	1

Loss Measurement State

Tx	Rx	Near End Loss Count	Far End Loss Count	Near End Loss Ratio	Far End Loss Ratio	Clear
0	0	0	0	0	0	<input type="checkbox"/>

Loss Measurement Availability

Enable	Interval	FLR Threshold	Maintenance
<input type="checkbox"/>	0	0	<input type="checkbox"/>

Loss Measurement Availability State

Near Availability Count	Far Availability Count	Near Unavailability Count	Far Unavailability Count	Near State	Far State
0	0	0	0	-	-

Loss Measurement High Loss Interval

Enable	FLR Threshold	Consecutive Interval
<input type="checkbox"/>	0	0

Loss Measurement High Loss Interval State

Near Count	Far Count	Near Consecutive Count	Far Consecutive Count
0	0	0	0

Loss Measurement Signal Degrade

Enable	TX Minimum	FLR Threshold	Bad Threshold	Good Threshold
<input type="checkbox"/>	0	0	0	0

Delay Measurement

Enable	Priority	Cast	Peer MEP	Ended	Tx Mode	Calc	Gap	Last-N	Unit	Synchronized	Counter Overflow Action
<input type="checkbox"/>	0	Multi	1	Single	Standardize	Flow	10	10	us	<input type="checkbox"/>	Keep

Delay Measurement State

	Tx	Rx	Rx Timeout	Rx Error	Av Delay Tot	Av Delay last N	Delay Min.	Delay Max.	Av Delay -Var Tot	Av Delay -Var last N	Delay -Var Min.	Delay -Var Max.	Overflow	Clear
One-way														
F-to-N	0	0	0	0	0	0	0	0	0	0	0	0	0	
N-to-F	0	0	0	0	0	0	0	0	0	0	0	0	0	
Two-way	0	0	0	0	0	0	0	0	0	0	0	0	0	<input type="checkbox"/>

Delay Measurement Bins

Measurement Bins for FD	Measurement Bins for IFDV	Measurement Threshold
3	3	5000

Delay Measurement Bins for FD

	bin0	bin1	bin2
One-way			
F-to-N	0	0	0
N-to-F	0	0	0
Two-way	0	0	0

Delay Measurement Bins for IFDV

	bin0	bin1	bin2
One-way			
F-to-N	0	0	0
N-to-F	0	0	0
Two-way	0	0	0

F-to-N :Far-end-to-near-end

N-to-F :Near-end-to-far-end

Loss Measurement

Loss Measurement (LM) offers a way for operators to determine the amount of frame loss in an Ethernet network (e.g., over an EVC). Specifically, LM is the ratio between undelivered OAM frames and the total number of OAM frames transmitted during a specific time interval. ITU-T Y.1731 defines two types of LM:

1. Single-Ended, where LM messages are transmitted to another MEP, which includes transmission and reception frame counts in its response message. Here, only the LM initiator is able to derive frame loss from the counters (since it does not include its local counters in the initial LM message); and
2. Dual-Ended. Continuity Check messages are used to carry frame transmission and reception counters. In contrast to the single-ended approach, this approach allows all MEPs inside a ME to derive frame loss, instead of only the initiating node.

Enable

Loss Measurement based on transmitting/receiving CCM or LMM/LMR PDU can be enabled/disabled - see 'Ended'. This is only valid with one Peer MEP configured.

Priority

The priority to be inserted as PCP bits in TAG (if any). In case of enable of Continuity Check and Loss Measurement both implemented on SW based CCM, 'Priority' must be the same for both.

Frame rate

Select the frame rate of CCM/LMM PDU. This is the inverse of transmission period described in Y.1731.

300f/sec selection is not valid.

100f/sec selection is not valid.

10f/sec selects ten frames-per-second as the frame rate of CCM/LMM PDU.

1f/sec selects one frame-per-second as the frame rate of CCM/LMM PDU.

6f/min selects six frames-per-minute as the frame rate of CCM/LMM PDU.

1f/min selects one frame-per-minute as the frame rate of CCM/LMM PDU.



If both Continuity Check and Loss Measurement are implemented and enabled on SW based CCM, the 'Frame Rate' must be the same for both.

Cast

Selection of CCM or LMM PDU transmitted **Uni**cast or **Multi**cast, where:

Uni: The unicast MAC will be taken from the 'Unicast Peer MAC' configuration. If both Continuity Check and dual ended Loss Measurement are implemented and enabled on SW based CCM, the 'Cast' setting must be the same on both.

Multi: CCM or LMM PDU transmitted multicast.

Ended

Single: Single-ended Loss Measurement implemented on LMM/LMR.

Dual: Dual-ended Loss Measurement implemented on SW based CCM.

FLR Interval

The interval in seconds where the calculated Frame Loss Ratio (FLR) is displayed (0-65535 seconds).
Loss Measurement State

Tx

The transmit count on which the LM is based.

Rx

the receive count on which the LM is based.

Near End Loss Count

The accumulated near end frame loss count - since the last 'Clear'.

Far End Loss Count

The accumulated far end frame loss count - since the last 'Clear'.

Near End Loss Ratio

The near end frame loss ratio calculated based on the near end frame loss count and far end frame transmitted - in the latest 'FLR Interval'. The result is given in percent.

Far End Loss Ratio

The far end frame loss ratio calculated based on the far end frame loss count and near end frame transmitted - in the latest 'FLR Interval'. The result is given in percent.

Clear

Check this checkbox and Save to clear the accumulated counters and restart ratio calculations.

Delay Measurement

Delay Measurement (DM) can be used for measuring delay in a Carrier Ethernet network. The unit of measurement is the round trip delay of a frame, measured from its first transmitted bit, until the reception of its last bit. Since a DM frame must be sent back to its originating node, LB messages are used.

Frame delay is the difference, in microseconds, between the time an ETH-DM frame is sent and received. (Frame delay variation - the difference between consecutive frame delay values - also called "frame jitter" - is a different parameter.)

Two types of DM can be identified:

One-way measurement: An initiating MEP includes a transmission timestamp in the Ethernet frame. The destination node will capture the frame reception timestamp, and compare both timestamps. As a consequence, the clocks of the sending and receiving nodes need to be synchronized; and

Two-way measurement: In contrast to the one-way measurement, this DM type does not require clock synchronization. The initiating node still includes a timestamp in the Ethernet frame. After the destination node performs a loopback on the frame, the initiating node will receive the frame again.

On reception, this node will capture the reception timestamp. Finally, the difference between the timestamps can be calculated.

Enable

Delay Measurement based on transmitting 1DM/DMM PDU can be enabled/disabled here. Delay Measurement based on receiving and handling 1DM/DMR PDU is always enabled.

Priority

The priority to be inserted as PCP bits in TAG (if any).

Cast

Selection of 1DM/DMM PDU transmitted **Unicast** or **Multicast**. The unicast MAC will be configured through 'Peer MEP'.

Peer MEP

This is only used if the 'Cast' is configured to 'Uni'. The 1DM/DMR unicast MAC will be taken from the 'Unicast Peer MAC' configuration of this peer. Frame delay measurement statistics are stored at this MEP for later retrieval.

Way

One-way: One-way Delay Measurement implemented on 1DM. The One-way delay measurement (1DM) PDU as defined in the ITU-T Y.1731 standard.

Two-way: Two-way Delay Measurement implemented on DMM/DMR. The DMM/DMR PDUs, defined in Y.1731, used to support on-demand two-way packet delay measurement or proactive two-way packet delay measurement.

Tx Mode

Standardize: Transmit 1DM/DMR per the ITU-T Y.1731 standard.

Proprietary: Transmit 1DM/DMR with follow-up packets using a proprietary method.

Calc

This is only used if the 'Way' is configured to Two-way.

Round trip: The frame delay calculated by the transmitting and receiving timestamps of initiators (Frame Delay = RxTimeb-TxTimeStampf).

Flow: The frame delay calculated by the transmitting and receiving timestamps of initiators and remotes. Frame Delay = (RxTimeb-TxTimeStampf)-(TxTimeStampb-RxTimeStampf).

Gap

The gap between transmitting 1DM/DMM PDU in 10 millisecond increments. The valid range is **10** to **65535** milliseconds.

Count

The number of last records to calculate. The range is **10** to **2000** records.

Unit

The time resolution. Select **us** or **ns**.

us: microseconds (uS or μS).

ns: nanoseconds (nS).

D2forD1

Enable to use DMM/DMR packet to calculate one-way DM. If the option is enabled, the following action will be taken. When DMR is received, two-way delay (roundtrip or flow) and both Near-end-to-far-end and Far-end-to-near-end one-way delay are calculated. When DMM or 1DM is received, only Far-end-to-near-end one-way delay is calculated. (Where **D2** indicates two-way delay, and **D1** indicates one-way delay measurement.)

Counter Overflow Action

The action the counter is to take when an overflow occurs.

Keep: Maintain the existing count when an overflow occurs.

Reset: Zero out the counter when an overflow occurs.

Delay Measurement State table

For one-way ETH-DM, only the receiver MEP (on the remote system) collects ETH-DM statistics.

Delay Measurement State

	Tx	Rx Timeout	Rx	Rx Error	Average Total	Average last N	Average Variation Total	Average Variation last N	Min.	Max.	Overflow	Clear
One-way												
F-to-N	0	0	0	0	0	0	0	0	0	0	0	
N-to-F	0	0	0	0	0	0	0	0	0	0	0	
Two-way	0	0	0	0	0	0	0	0	0	0	0	<input type="checkbox"/>

One-way DM performs one-way ETH-DM, which is based on the difference between the time at which the initiator MEP sends a one-way ETH-DM delay measurement request (1DM) frame and the time at which the receiver MEP receives the frame.

F-to-N (Far-end-to-near-end) One-way Delay

The one-way delay is from remote devices to local devices. The conditions to calculate this delay are:

- 1) 1DM received.
- 2) DMM received with D2forD1 enabled.
- 3) DMR received with D2forD1 enabled.

	Tx	Rx Timeout	Rx
One-way			
F-to-N	0	0	0
N-to-F	0	0	0
Two-way	0	0	0

F-to-N :Far-end-to-near-end

N-to-F :Near-end-to-far-end

N-to-F (Near-end-to-far-end) One-way Delay

The one-way delay is from local devices to remote devices. The only case to calculate this delay is:

- 1) DMR received with D2forD1 enabled.

Two-way Delay

Performs two-way ETH-DM, which is based on the difference between the time at which the initiator MEP sends a two-way ETH-DM delay measurement message (DMM) frame and the time at which

the initiator MEP receives an associated two-way ETH-DM delay measurement reply (DMR) frame from the responder MEP, subtracting the time elapsed at the responder MEP.

Tx

Displays the accumulated transmit count - since the last 'Clear'.

Rx Timeout

The accumulated receive timeout count for two-way only - since the last 'Clear'.

Rx

The accumulated receive count - since the last 'Clear'.

Rx Error

The accumulated receive error count - since the last 'Clear'. The frame delay is larger than 1 second (timeout).

Average Total

The average delay - since the last 'Clear'. The unit is microseconds.

Average last N

The average delay of the last n packets - since the last 'Clear'. The unit is microseconds.

Average Variation Total

The average delay variation - since the last 'Clear'. The unit is microseconds.

Average Variation last N

The average delay variation of the last n packets - since the last 'Clear'. The unit is microseconds.

Min.

The minimum delay since the last 'Clear'. The unit is microseconds.

Max.

The maximum delay since the last 'Clear'. The unit is microseconds.

Overflow

The number of counter overflows since the last 'Clear'.

Clear

Check this checkbox and the next Save will clear the accumulated counters.

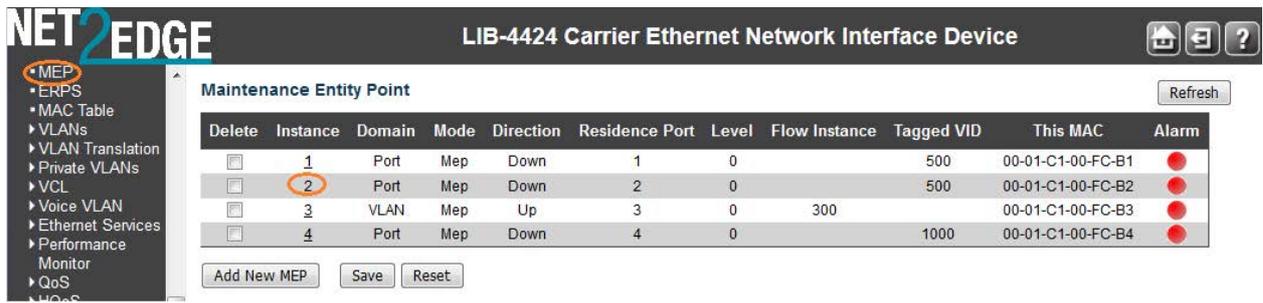
Buttons

Save: Click to save changes.

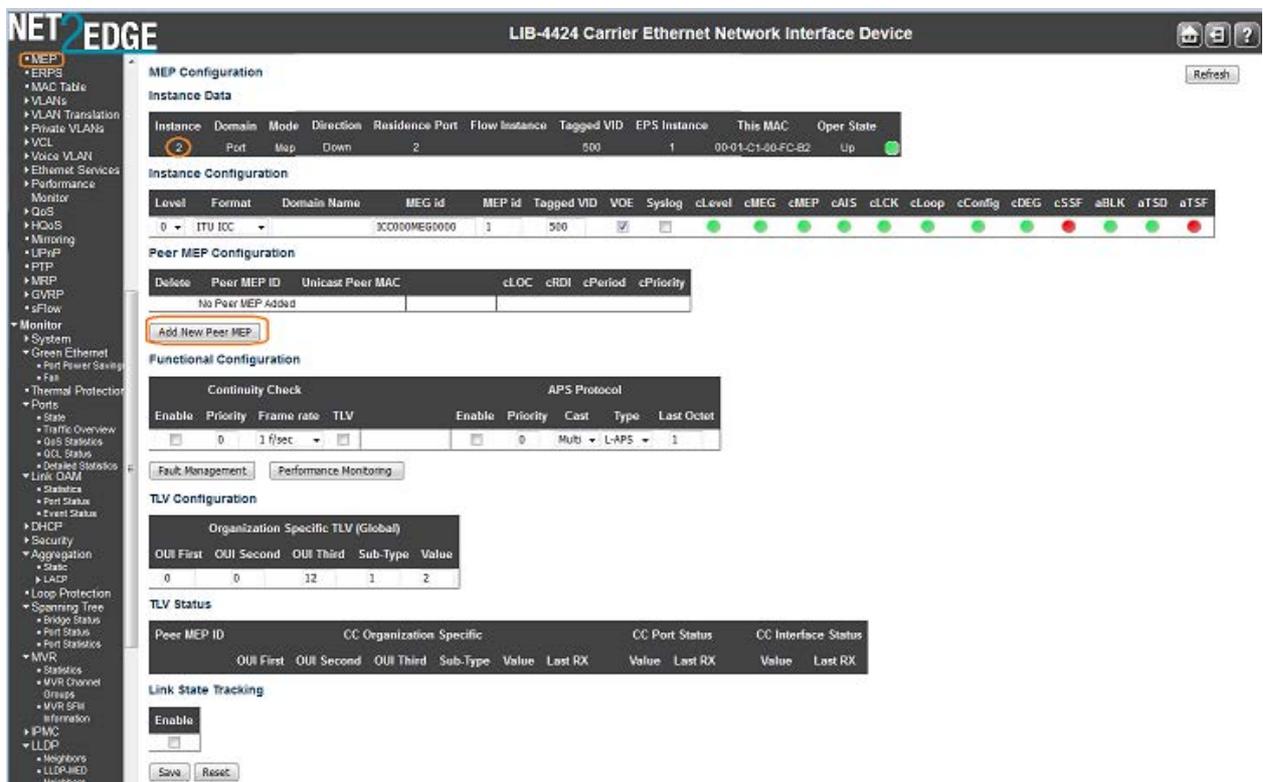
Reset: Click to undo any changes made locally and revert to previously saved values.

Add a New Peer MEP Procedure

Navigate to the **Configuration > MEP** menu path. See “MEP Instance Configuration” on page 162.

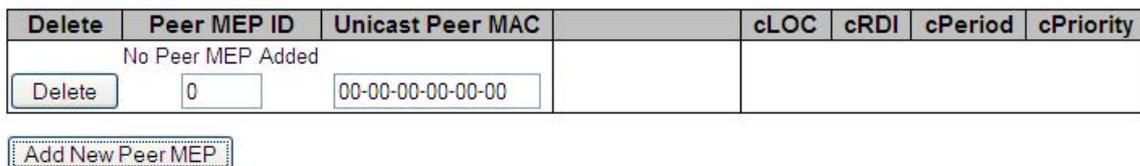


Select a MEP Instance to configure (e.g., click on the 2 in the **Instance** column above). The MEP Instance 2 config page displays.

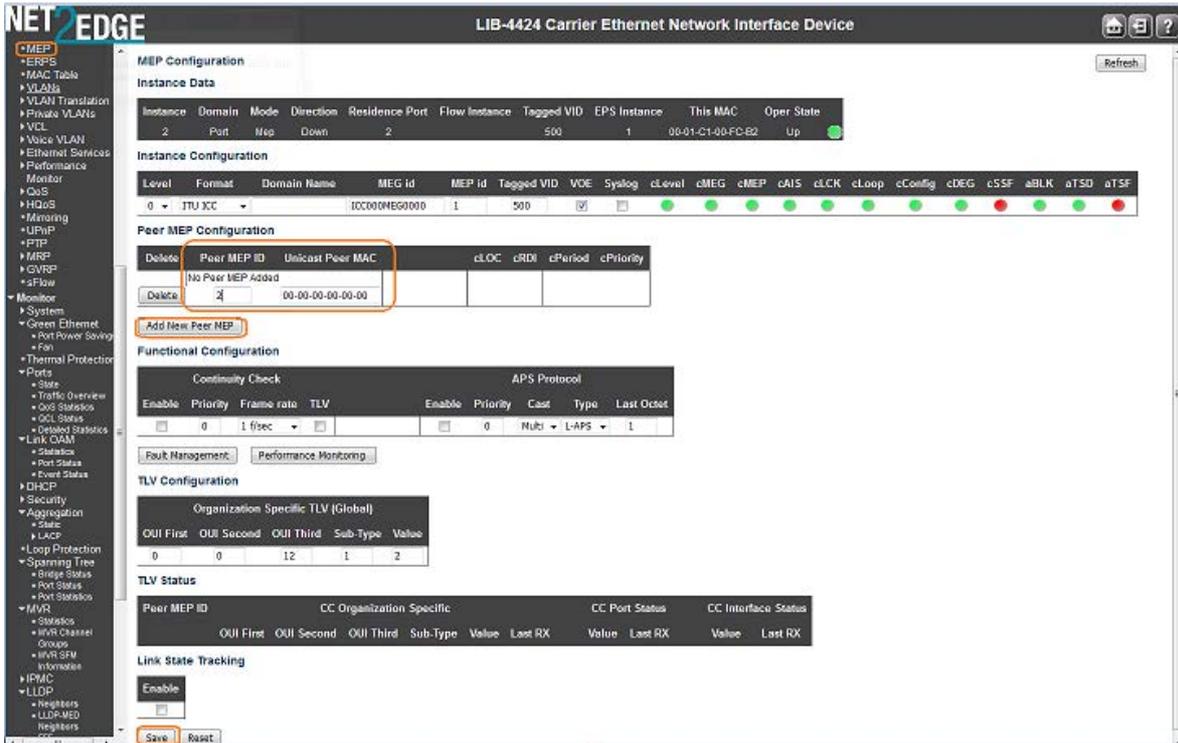


In the “Peer MEP Configuration” section, click the **Add New Peer MEP** button. An entry table displays with the message “No Peer MEP Added.”

Peer MEP Configuration

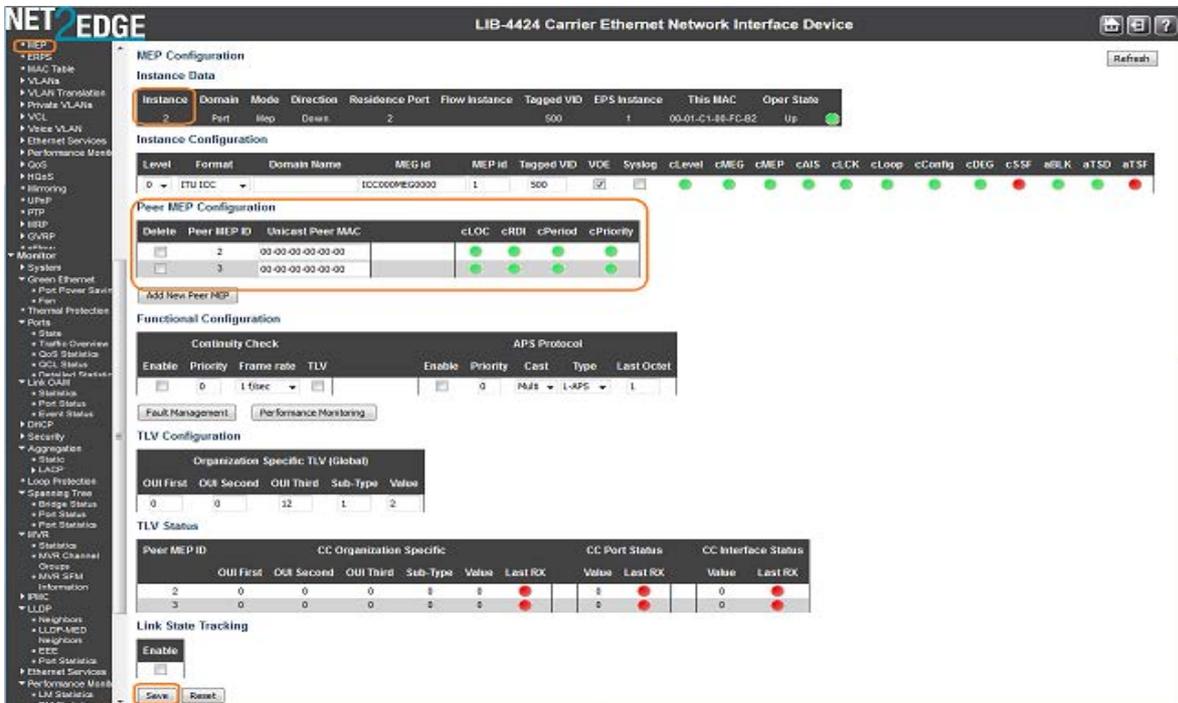


Enter a **unique Peer MEP ID** for the new peer MEP (you cannot enter duplicate MEP IDs).
 Enter a **Unicast Peer MAC** address for the new peer MEP (this must be a unicast address).



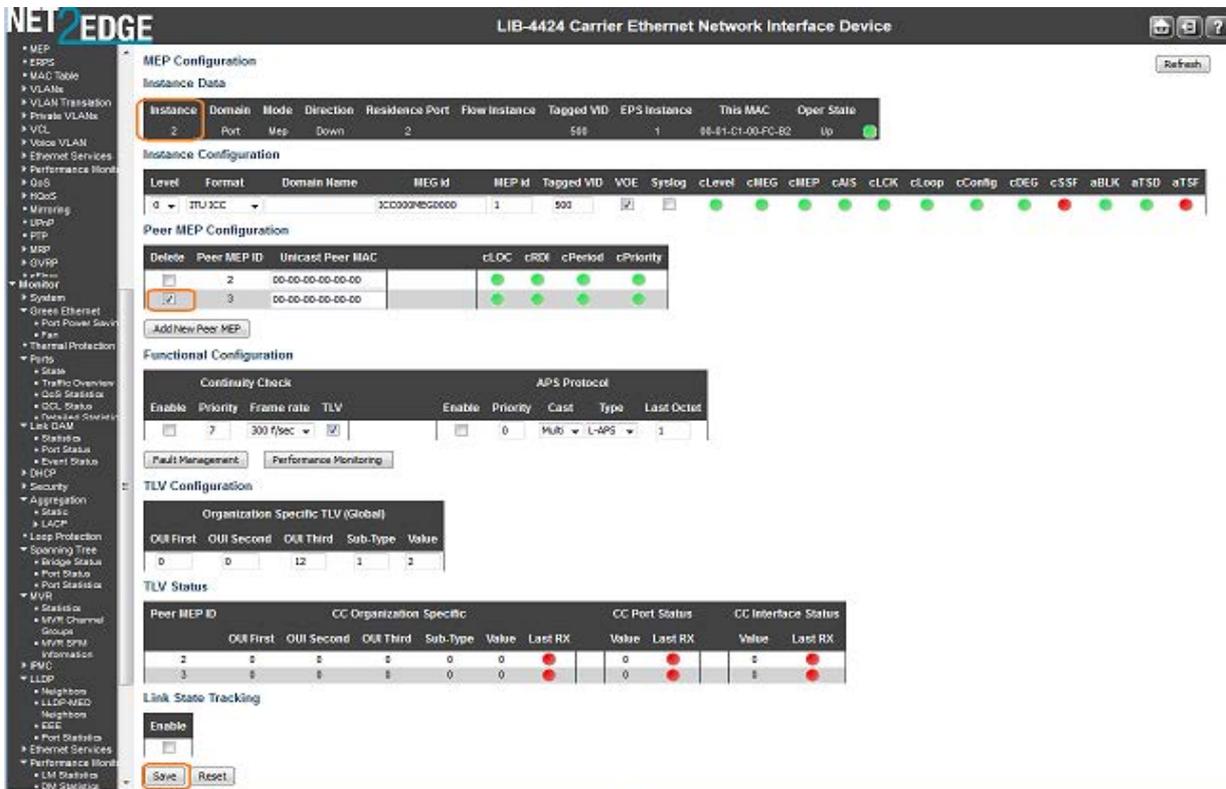
Click the **Save** button. Note that only one peer MEP can be added per save operation. Repeat steps 2-6 for each new peer MEP to be added.

Verify your new MEP peer configurations (e.g., Peer MEP IDs 0 and 1 shown below).

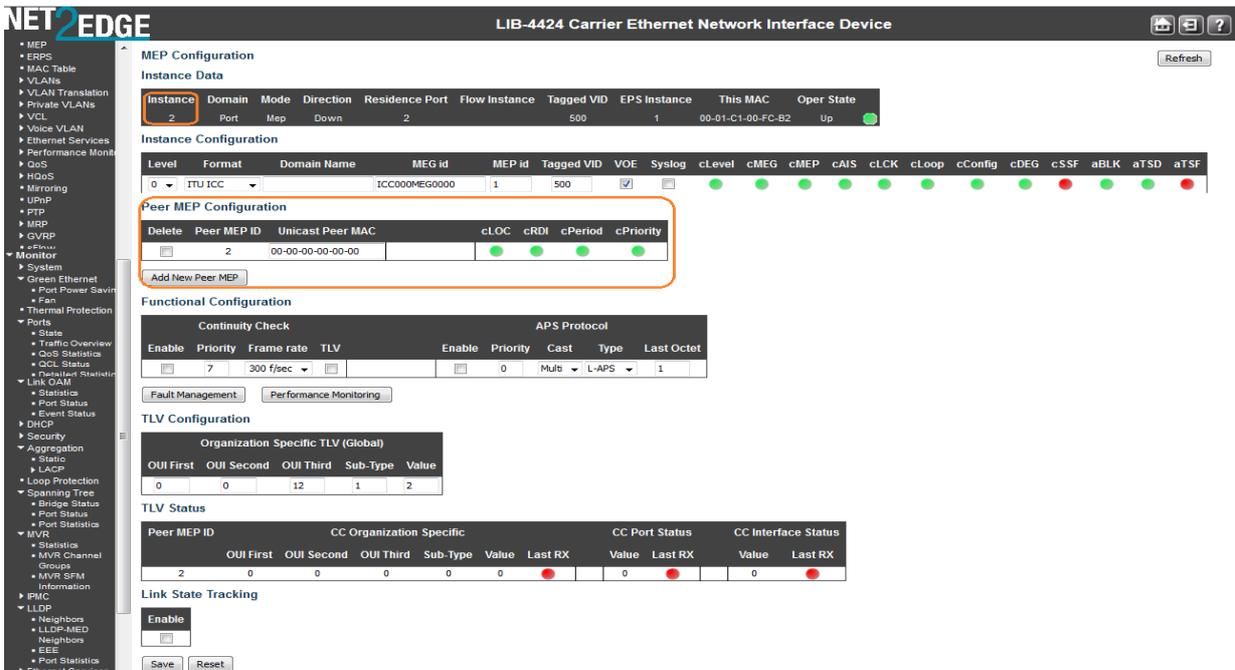


Delete a Peer MEP Procedure

To delete an existing Peer MEP, check its checkbox in the **Delete** column and then click the **Save** button.

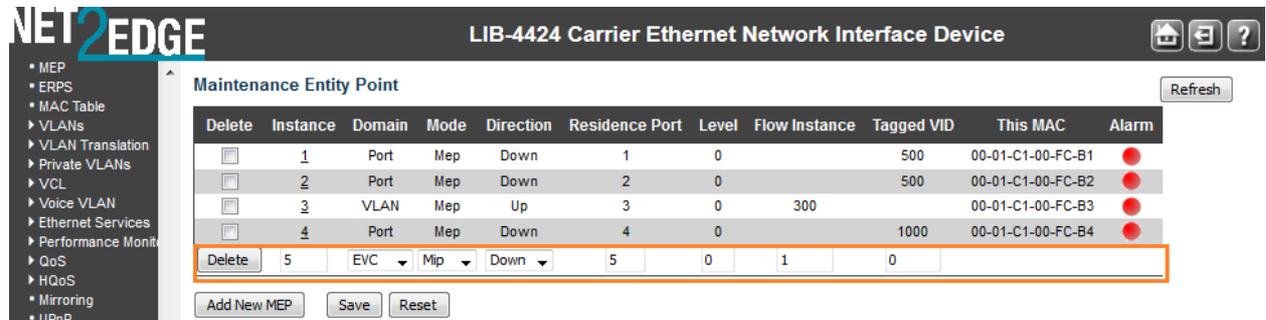


Verify that the Peer MEP was deleted from the table. Click the **Refresh** button if necessary.



Add a New MIP Procedure

You can create a MIP on an EVC, but not on a Port. You select EVC as the Domain and select Up as the Direction, and then select the Level. There is no option to configure a MEG ID, etc. Navigate to the **Configuration > Spanning Tree > CIST Ports** menu path and disable STP. Navigate to the **Configuration > VLAN** menu path and configure VLANs as required. Navigate to the **Configuration > MEP** menu path and click the **Add New MEP** button.



Enter an **Instance** number.
 At the **Domain** dropdown, select **EVC**.
 At the **Mode** dropdown, select **MIP**.
 At the **Direction** dropdown, select **Up**.
 At the **Residence Port**, **Level**, **Flow Instance**, and **Tagged VID** dropdowns enter valid values (see above).
 Click the **Save** button when done.

Difference in IEEE 802.1ag and ITU-T 1731 SOAM Test Terms

The SOAM MEP tests vary slightly between the IEEE 802.1ag and ITU-T 1731 standards, as noted below:

Table 2: SOAM Test Terms - IEEE 802.1ag vs. ITU-T 1731

Function	Feature	802.1ag	Y.1731
Connectivity Fault Management	Continuity Check	Yes	Yes
	Loopback	Yes	Yes
	Link Trace	Yes	Yes
	AIS	Not applicable	Yes
	RDI	Yes	Yes
Performance Monitoring	ETH-TST	Not applicable	Yes
	Loss Measurement	Not applicable	Yes
	Delay Measurement and Delay Variation Measurement	Not applicable	Yes

ERPS Configuration

LIB-44xx ERPS (Ethernet Ring Protection Switching) is configured from the **Configuration > ERPS** menu path. An ERP instance is an entity that is responsible for the protection of a subset of the VLANs that transport traffic over the physical Ethernet ring. Each ERP instance is independent of other ERP instances that may be configured on the physical Ethernet ring.

The LIB-44xx implements the ITU G.8032 standard for ERPS, which uses the APS automatic protection protocol for protection in ring and interconnected ring topology. The LIB-44xx supports G.8032v1 in a single ring topology and G.8032v2 in multiple rings/ladder topologies.

The G.8032 protocol is based on the SONET protection capability and requires sub 50ms switchover in case of failure. In a ring topology, loops are prevented by blocking one of the links in the ring (the “Ring Protection Link”).

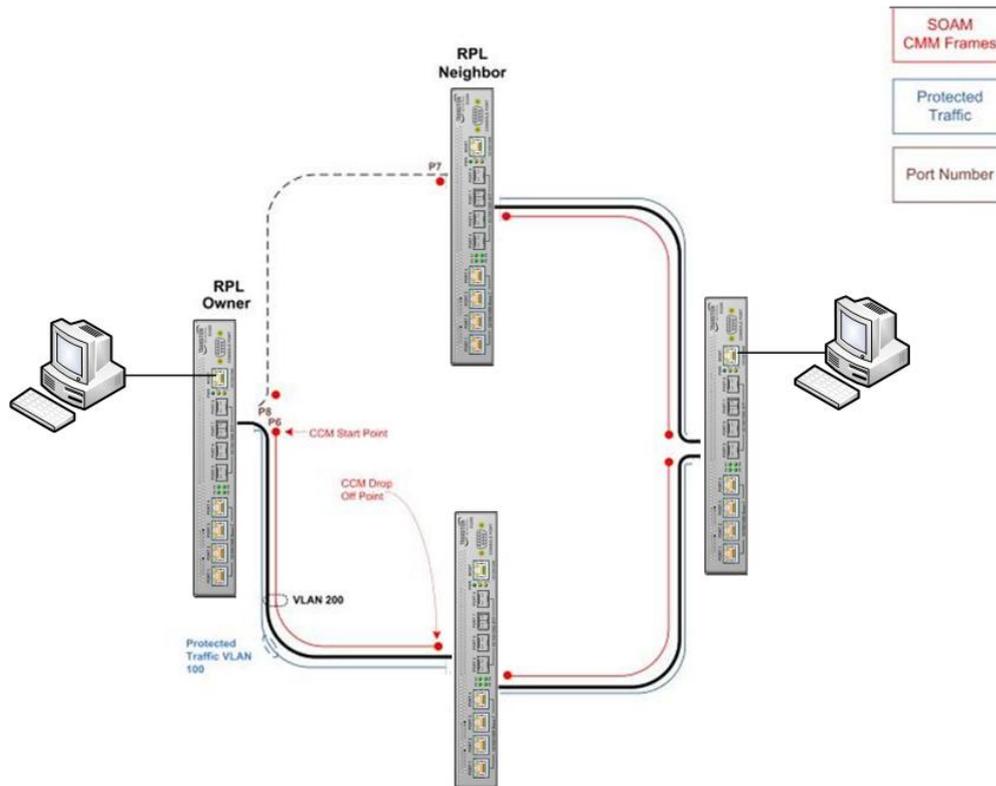


Figure 4. ERPS Example

Ring protection is based on loop avoidance. This is achieved by guaranteeing that at any time traffic may flow on all but one of the ring links. This principle derives the following rule for the protection control protocol: “*once a ring port has been blocked, it may be unblocked only if it is known that there remains at least one other blocked ring port in the Ethernet ring.*” This rule is used as the basis to control all actions of traffic channel unblocking in the Ethernet ring, as well as to define the information necessary to distribute between all Ethernet ring nodes. The protection algorithm is based on the transmission of local switch requests and local status to all Ethernet ring nodes via the R-APS specific information.

The designated node that does the blocking is called the RPL ‘owner’ and the node at the other end of the RPL is the RPL ‘neighbour’. The nodes on the ring use the R-APS protocol to co-ordinate activating protection on the ring.

ERPS Sample Setups 1

The figures below are provided to help explain interconnected rings, interconnection nodes, and ring links in terms of ERPS user setup.

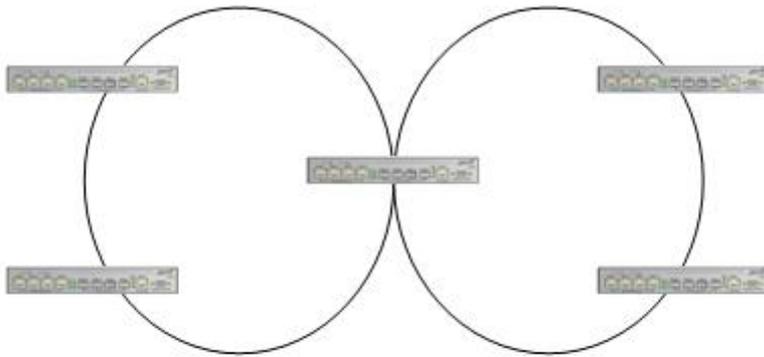


Figure 5. Interconnected Ethernet Rings via an Interconnection Node

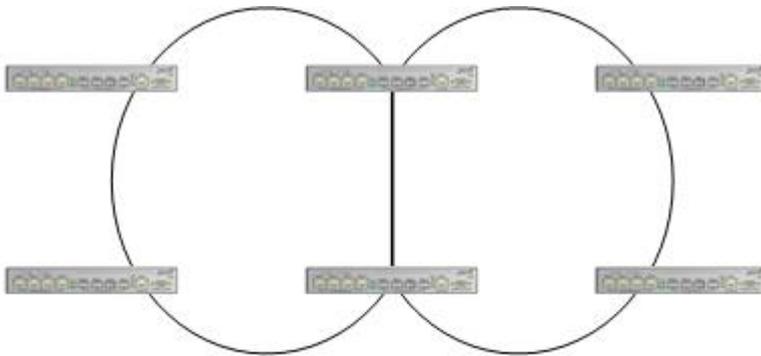


Figure 6. Interconnected Ethernet Rings via Dual Nodes with a Ring Link

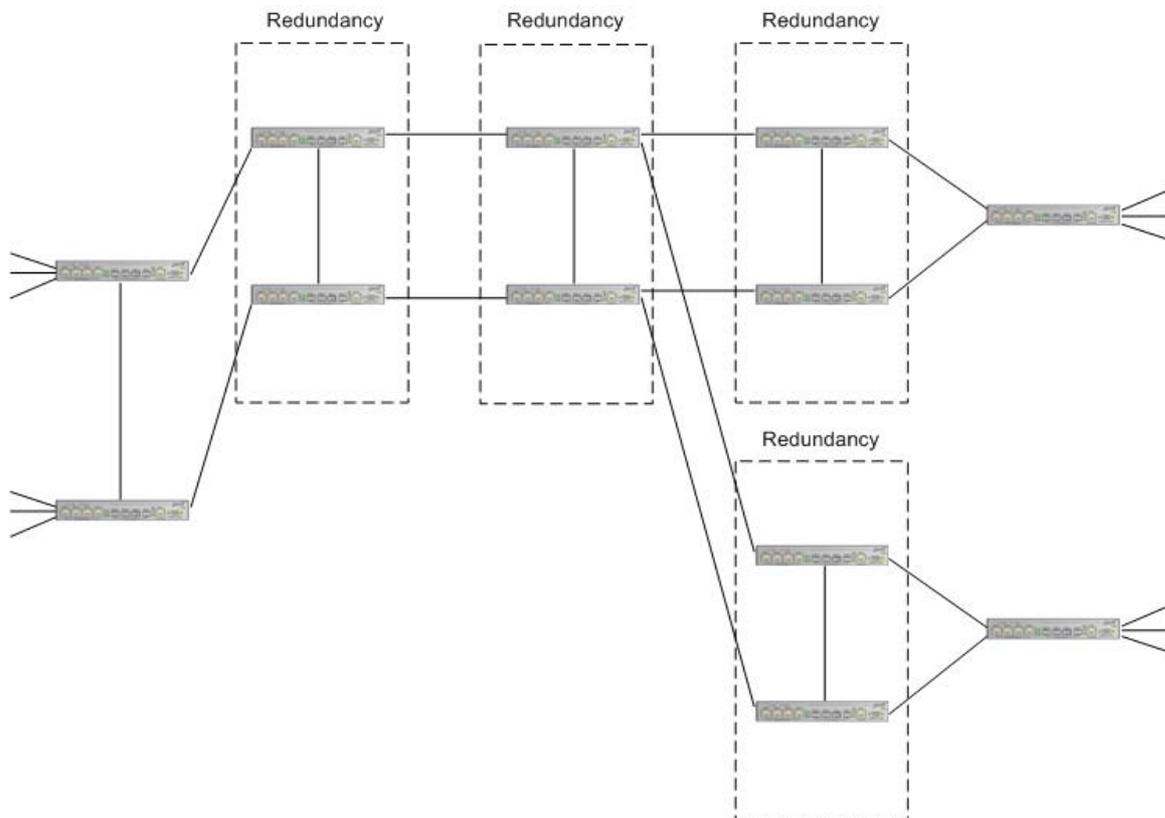


Figure 7. Ethernet Multi-Ring / Ladder Network

ERPS Sample Setups 2

The figures below are provided to help explain sub-rings, intermediate nodes, interconnection nodes, major rings, and multiple sub-rings in terms of RPL user setup.

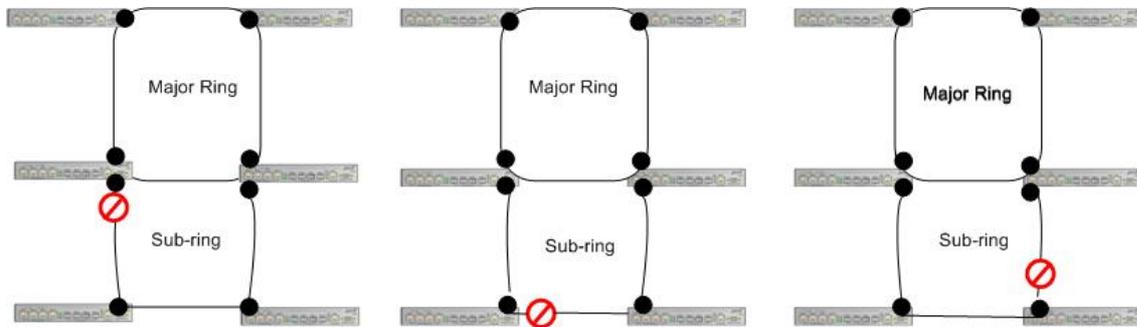


Figure 8. Location of RPL for a Sub-ring

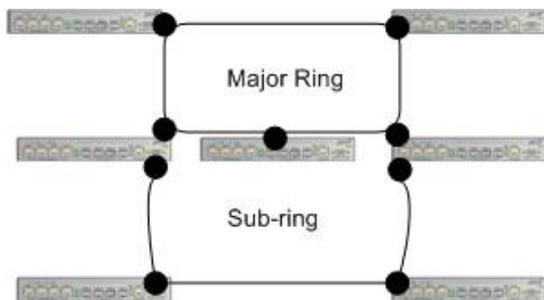


Figure 9. Intermediate Nodes between Interconnection Nodes

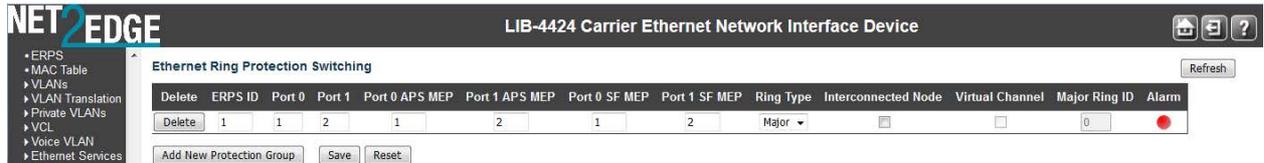


Figure 10. Multiple Sub-rings connected to a Major Ring

When you select the **Configuration > ERPS** menu path, the default page displays.



At the default Ethernet Ring Protection Switching page, click the **Add New Protection Group** button to display the ERPS table and entry parameters.



The ERPS parameters are explained below:

Delete

This checkbox is checked to mark an ERPS for deletion in the next Save operation.

ERPS ID (Protection Group ID)

The ID of the new Protection group. Enter an ID value of **1-64**. You can click on the ID of an existing Protection group to enter its configuration page (described later in this section).

Port 0

This will create a Port 0 of the switch in the ring. Assign an integer value of **1-4**. The Port 0 and Port 1 cannot be the same.

Port 1

This will create "Port 1" of the switch in the Ring. As interconnected sub-ring will have only one ring port, "Port 1" is configured as **"0"** for interconnected sub-ring. A **"0"** in this field indicates that no "Port 1" is associated with this instance. Assign an integer value of **1-4**. The Port 0 and Port 1 cannot be the same.

Port 0 APS MEP

This is the Port 0 APS PDU handling MEP. Assign an integer value of **1-32**. The Port 0 APS MEP and Port 1 APS MEP cannot be the same. Note that ERPS MEPs are referenced by Instance number instead of MEP ID number. The MEP Instance number must be entered, which may or may not match the MEP ID.

Port 1 APS MEP

The Port 1 APS PDU handling MEP. As only one APS MEP is associated with interconnected sub-ring without virtual channel, it is configured as **"0"** for such ring instances. **"0"** in this field indicates that no Port 1 APS MEP is associated with this instance. Assign an integer value of **1-32**. The Port 0 APS MEP and Port 1 APS MEP cannot be the same. Note that ERPS MEPs are referenced by Instance number instead of MEP ID number. The MEP Instance number must be entered, which may or may not match the MEP ID.

Port 0 SF MEP

This is the Port 0 Signal Fail reporting MEP. Assign an integer value of **1-32**. Port 0 SF MEP and Port 1 SF MEP cannot be the same. Note that ERPS MEPs are referenced by Instance number instead of MEP ID number. The MEP Instance number must be entered, which may or may not match the MEP ID.

Port 1 SF MEP

This is the Port 1 Signal Fail reporting MEP. As only one SF MEP is associated with interconnected sub-ring without virtual channel, it is configured as "0" for such ring instances. A "0" in this field indicates that no Port 1 SF MEP is associated with this instance. Assign an integer value of **1-32**. Port 0 SF MEP and Port 1 SF MEP cannot be the same. Note that ERPS MEPs are referenced by Instance number instead of MEP ID number. The MEP Instance number must be entered, which may or may not match the MEP ID.

Ring Type

Select the type of Protection ring. It can be either **Major** ring or **Sub**-ring.

Major ring: the Ethernet ring that is connected on two ports to an interconnection node.

Sub-ring: an Ethernet ring which is connected to one or more other Ethernet rings or networks through the use of a pair of interconnection nodes. On their own, the sub-ring links do not form a closed loop. A closed connection of traffic may be formed by the sub-ring links and one or more links that are controlled by other Ethernet ring or network, between interconnection nodes.

Interconnected Node

Interconnected Node indicates that the ring instance is interconnected. Click on the checkbox to configure this.

"**Yes**" indicates it is an interconnected node for this instance.

"**No**" indicates that the configured instance is not interconnected.

Virtual Channel

Sub-rings can either have virtual channel or not on the interconnected node. This is configured using this "Virtual Channel" checkbox.

"**Yes**" indicates it is a sub-ring with virtual channel.

"**No**" indicates, sub-ring doesn't have virtual channel. (Sub-ring with or without R-APS virtual channel.)

Sub-ring with R-APS virtual channel: In this option, a virtual channel to tunnel R-APS messages from one interconnection node to the other interconnection node is established.

Sub-ring without R-APS virtual channel: In this option, the R-APS channel is terminated at the interconnection nodes and its R-APS messages are not tunnelled between the interconnection nodes.

The Ring Interconnection options are shown in Rec. ITU-T G.8032/Y.1344 (03/2010).

R-APS messages are transmitted with the request/state and status information defined by the R-APS request process. The R-APS messages are transported via an R-APS specific VLAN. If the R-APS information to be transmitted has been changed, a burst of three R-APS messages is transmitted as quickly as possible, to ensure the fastest protection switching possible. For messages other than an 'event' message, the R-APS message continues to be transmitted, after the first three messages are transmitted, with a frequency of one message every five seconds. Typically, R-APS messages are transmitted on both ring ports.

Major Ring ID

This is the Major ring group ID for the interconnected sub-ring. It is used to send topology change updates on major ring. If ring is Major, this value is same as the protection group ID of this ring.

Alarm

There is an active alarm on the ERPS. A Green dot = Up, a red dot = Down.

Buttons

Add new Protection Group: Click to add a new Protection group entry.

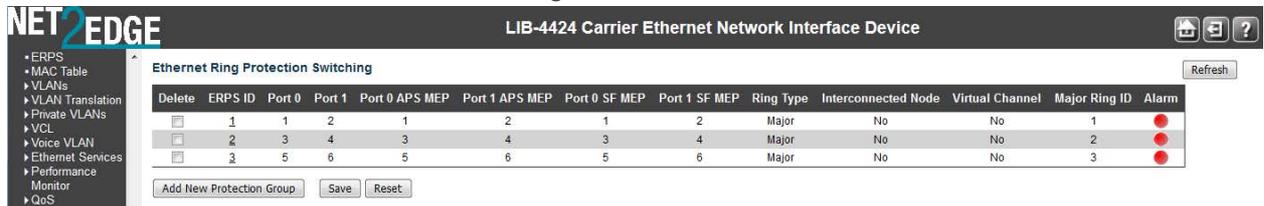
Refresh: Click to refresh the page immediately.

Save: Click to save changes. Only one ERPS can be added per save operation.

Reset: Click to undo any changes made locally and revert to previously saved values.

Example

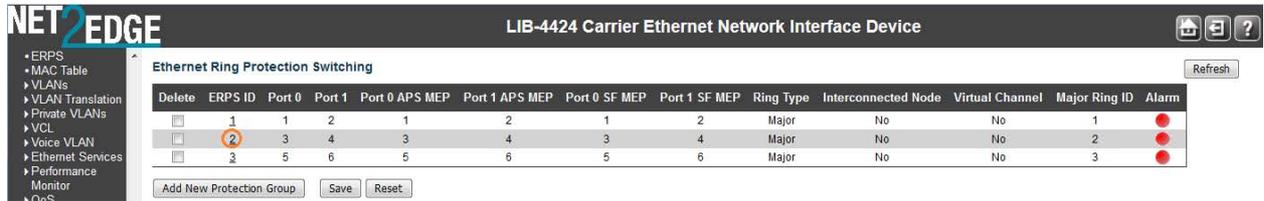
The screen below shows four ERPS IDs configured.



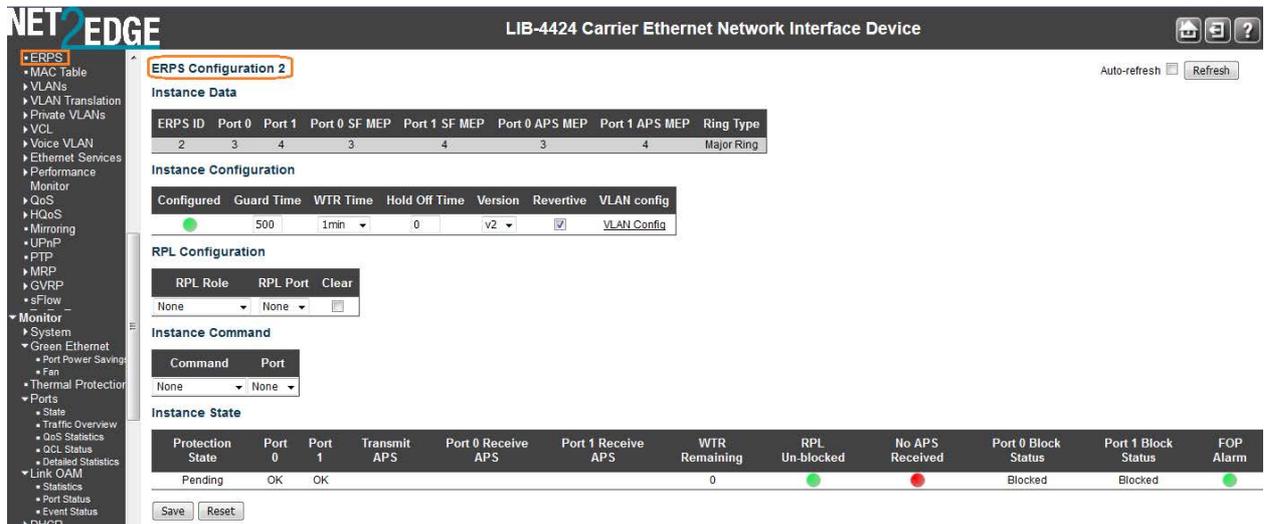
When you click on the linked ERPS ID of an existing Protection group, its ERPS Configuration page displays, as discussed below:

ERPS Configuration Page

When you click on the ERPS ID of an existing Protection group, its configuration page displays.



For example, if you click on ERPS ID 2 on the screen above, the ERPS ID 2 configuration displays:



This screen lets you configure the ERPS Instance Data, Instance Configuration, RPL Configuration, Sub-Ring Configuration, and Instance Command, and view the ERPS Instance State parameters. These parameters are explained below:

ERPS Instance Data

ERPS ID

The ID of the new Protection group. Enter an ID value of **1-64**. Click on the ID of an existing Protection group to enter its configuration page.

Port 0

This will create a Port 0 of the switch in the ring. Assign an integer value of **1-4**. The Port 0 and Port 1 cannot be the same.

Port 1

This will create "Port 1" of the switch in the Ring. As interconnected sub-ring will have only one ring port, "Port 1" is configured as "0" for interconnected sub-ring. "0" in this field indicates that no "Port 1" is associated with this instance. Assign an integer value of **1-4**. The 'Port 0' and 'Port 1' entries cannot be the same.

Port 0 SF MEP

This is the Port 0 Signal Fail reporting MEP. Assign an integer value of **1-32**. Port 0 SF MEP and Port 1 SF MEP cannot be the same. Note that ERPS MEPs are referenced by Instance number instead of MEP ID number. The MEP Instance number must be entered, which may or may not match the MEP ID.

Port 1 SF MEP

This is the Port 1 Signal Fail reporting MEP. As only one SF MEP is associated with interconnected sub-ring without virtual channel, it is configured as "0" for such ring instances. A "0" in this field indicates that no Port 1 SF MEP is associated with this instance. Assign an integer value of **1-32**. Port 0 SF MEP and Port 1 SF MEP cannot be the same. Note that ERPS MEPs are referenced by Instance number instead of MEP ID number. The MEP Instance number must be entered, which may or may not match the MEP ID.

Port 0 APS MEP

This is the Port 0 APS PDU handling MEP. Assign an integer value of **1-32**. The Port 0 APS MEP and Port 1 APS MEP cannot be the same. Note that ERPS MEPs are referenced by Instance number instead of MEP ID number. The MEP Instance number must be entered, which may or may not match the MEP ID.

Port 1 APS MEP

The Port 1 APS PDU handling MEP. As only one APS MEP is associated with interconnected sub-ring without virtual channel, it is configured as "0" for such ring instances. "0" in this field indicates that no Port 1 APS MEP is associated with this instance. Assign an integer value of **1-32**. The Port 0 APS MEP and Port 1 APS MEP cannot be the same. Note that ERPS MEPs are referenced by Instance number instead of MEP ID number. The MEP Instance number must be entered, which may or may not match the MEP ID.

Ring Type

Type of Protection ring. It can be either **Major** ring or **Sub**-ring.

Major ring: the Ethernet ring that is connected on two ports to an interconnection node.

Sub-ring: an Ethernet ring which is connected to one or more other Ethernet rings or networks through the use of a pair of interconnection nodes. On their own, the sub-ring links do not form a closed loop. A closed connection of traffic may be formed by the sub-ring links and one or more links, that are controlled by other Ethernet ring(s) or network(s), between interconnection nodes.

ERPS Instance Configuration

Configured

Displays a green LED (●) for Up or a red LED (●) for Down.

Red: This ERPS is only created and has not yet been configured and is not active.

Green: This ERPS is configured and is active.

Guard Time

Enter the Guard timeout value to be used to prevent ring nodes from receiving outdated R-APS messages. The period of the guard timer can be configured in 10 ms steps between **10 ms** and **2 seconds**, with a default value of **500 ms**. The guard timer is used to prevent Ethernet ring nodes from acting on outdated R-APS messages, and prevents the possibility of forming a closed loop. The guard timer is activated whenever an Ethernet ring node receives an indication that a local switching request has cleared (i.e., Local Clear SF, Clear). This guard timer period should be greater than the maximum expected forwarding delay in which an R-APS message traverses the entire ring. The longer the period of the guard timer, the longer an Ethernet ring node is unaware of relevant new or existing requests transmitted from other Ethernet ring nodes, and therefore unable to react to them. A guard timer is used in every Ethernet ring node. Once a guard timer is started, it expires by itself. When the guard timer is not running, the R-APS request/state and status information is forwarded unchanged.

WTR Time

The Wait To Restore timing value to be used in revertive switching. You can set the WTR period to **1** minute or from **5-12** minutes in 1 minute steps. The default value is **5** minutes. In revertive mode, the wait to restore (WTR) timer is used to prevent frequent operation of the protection switching due to intermittent signal failure defects.

Hold Off Time

The timing value to be used to make persistent check on Signal Fail (SF) before switching. The Hold off timer valid range is **0** to **10** seconds in steps of 100 ms. The default is **0**.

The holdoff timer is used to allow, for example, a server layer protection switch to have a chance to fix the problem before switching at a client layer, in order to coordinate timing of protection switches at multiple layers. When a new defect or more severe defect occurs (new SF), this event is not reported immediately to protection switching if the provisioned holdoff timer value is non-zero. Instead, the holdoff timer is started. When the holdoff timer expires, the trail that started the timer is checked as to whether a defect still exists. If one does exist, that defect is reported to protection switching. The reported defect need not be the same one that started the timer.

Version

Select **v1** or **v2** as the ERPS version to be used. For fields such as Version, OpCode, Flags, and End TLV, the values used are as defined in ITU-T Y.1731 (Version 0x01 is transmitted per the current version of this Recommendation at the time of this publication.) G.8032v1 supported a single ring topology and G.8032v2 supports multiple rings/ladder topology.

v1: G.8032 v1 supported a single ring topology. The v1 protocol is robust enough to work for

unidirectional failure and multiple link failure scenarios in a ring topology. It allows mechanism to force switch (FS) or manual switch (MS) to take care of field maintenance scenario.

v2: G.8032 v2 supports multiple rings/ladder topology. The v2 protocol also introduced other features such as Revertive/ Non-revertive mode after condition, that is causing the switch, is cleared, Administrative commands - Forced Switch (FS), Manual Switch (MS) for blocking a particular ring port, Flush FDB (Filtering database), and support of multiple ERP instances on a single ring.

Revertive

Check the checkbox for Revertive mode operation. Uncheck the checkbox for NonRevertive mode operation. An Ethernet ring node that has one or more ring ports in an SF condition, upon detection of clearance of the SF condition, keeps at least one of these ring ports blocked for the traffic channel and for the R-APS channel, until the RPL is blocked as a result of Ethernet ring protection reversion, or until there is another higher priority request (e.g., an SF condition) in the Ethernet ring.

- Revertive operation:** When all ring links and Ethernet ring nodes have recovered and no external requests are active, reversion is the action to be taken.
- Non-revertive operation:** the Ethernet ring does not automatically revert when all ring links and Ethernet ring nodes have recovered and no external requests are active.

Both revertive and non-revertive handling are discussed in Rec. ITU-T G.8032/Y.1344 (03/2010). Protection switching on a manual switch request is completed when the specified actions are performed by each Ethernet ring node. At this point, the conditions are created to allow the traffic flows to be steered around the Ethernet ring.

VLAN Config

You can click the [VLAN Config](#) hyperlink to display the ERPS VLAN Configuration for the port. See below for description.

RPL Configuration

The ring protection link is the ring link that under normal conditions (i.e., without any failure or request) is blocked (at one or both ends) for traffic channel, to prevent the formation of loops.

RPL Role

Select either '**RPL_Owner**' or '**RPL_Neighbour**', where:

RPL Neighbour node, when configured, is an Ethernet ring node adjacent to the RPL that is responsible for blocking its end of the RPL under normal conditions (i.e., the ring is established and no requests are present in the ring) in addition to the block by the RPL owner node. However, the RPL Neighbour is not responsible for activating the reversion behaviour.

RPL Owner node is an Ethernet ring node adjacent to the RPL that is responsible for blocking its end of the RPL under normal conditions (i.e., the ring is established and no requests are present in the ring).

It is also responsible for activating reversion behaviour from protected or manual switch/forced switch (MS/FS) conditions.

RPL Port

This dropdown lets you select the east port or west port as the RPL block.

None: Nothing selected as the RPL block.

Port0: This selects the East port of the LIB-44xx in the ring as the RPL block.

Port1: This selects the West port of the LIB-44xx in the ring.

Clear

Clear: If the owner must be changed; check the Clear checkbox to clear the RPL owner for that ERPS ring.

Sub-Ring Configuration

Displays only for a “Sub” Ring type instance.

Ring Type

Displays “Sub” Ring or “Major” as the ring type for this instance.

Topology Change

Check this checkbox to cause topology changes in the sub-ring to be propagated to the Major ring. (This field only displays for a Ring Type of ‘Sub-Ring’. Topology change propagation, when enabled, sends a *Topology_Change* signal when a flush FDB action is triggered by the ERP Control Process of a Sub-Ring’s ERP Instance. The *Topology_Change* signal is disabled after a period of 10 ms.

Instance Command

Command

A port (e.g., Port0 or Port1) can be administratively configured to be in either **Manual** switch or **Forced** switch state or **None** (neither Forced or Manual) from the Command dropdown, or you can **Clear** the active local administrative selection.

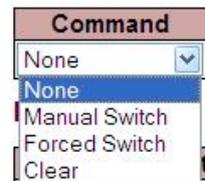
None (neither Forced or Manual) from the Command dropdown (the default).

Manual Switch: In the absence of a failure or FS, a Manual Switch selection forces a block on the ring port where the command is issued.

Forced Switch: A Forced Switch (FS) selection forces a block on the ring port where the command is issued.

Clear: The Clear entry clears the active local administrative selection (e.g., Forced Switch or Manual Switch state). The Clear command is used for these operations:

- Clearing an active local administrative command (e.g., forced switch or manual switch),
- Triggering reversion before the WTR or WTB timer expires in case of revertive operation, and
- Triggering reversion in case of non-revertive operation.



Port

The port (**None**, **Port0** or **Port1**) to be administratively configured.

Instance State

Protection State

The current ERPS state according to State Transition Tables in G.8032.

Pending: The state is in the process of changing; you have selected a change but not yet ‘Saved’ the change.

Protected: ERPS protected mode is enabled for this instance.

None: No protected mode is enabled for this instance.

Idle: Protected mode is idle for this instance.

Forced Switch: this instance is in forced switch mode. A Forced Switch (FS) selection forces a block on the ring port where the command is issued.

Manual Switch: this instance is in manual switch mode. In the absence of a failure or FS, a Manual Switch selection forces a block on the ring port where the command is issued.

See “*Ethernet linear protection switching - Recommendation ITU-T G.8031/Y.1342, Annex A*” – State transition tables of protection switching Tables A.1 - A.6. These tables provide protection switching state transition information for various protection switching configurations (although Annex A does not form an integral part of the Recommendation). The states include, but are not necessarily limited to No request (NR), Lockout (LO), Forced switch (FS), Signal fail (W) SF, Signal fail (P) SF-P, Manual switch MS, Wait to restore WTR, Exercise EXER, Reverse request RR). Annex A notes that any other global or local request which is not described in the state transition tables does not trigger any state transition.

Port 0

OK: State of East port is ok.

SF: State of East port is Signal Fail.

Port 1

OK: State of West port is ok.

SF: State of West port is Signal Fail.

Transmit APS

The transmitted APS according to the State Transition Tables in G.8032. Signal Fail (SF) is declared when ETH trail signal fail condition is detected. No Request (NR) is declared when there are no outstanding conditions (e.g., SF, SF DNF BPR1, NR BPR0, etc.) on the node. See Rec. ITU-T G.8031/Y.1342 for details.

Port 0 Receive APS

The received APS for Port 0 according to State Transition Tables in G.8032.

Port 1 Receive APS

The received APS for Port 1 according to State Transition Tables in G.8032.

WTR Remaining

The remaining WTR (Wait to Restore) timeout in milliseconds.

RPL Un-blocked

APS is received on the working flow. Displays a green LED (●) for Up or a red LED (●) for Down.

No APS Received

RAPS PDU is not received from the other end. Displays a green LED (●) for Up or a red LED (●) for Down.

Port 0 Block Status

Block status for Port 0 (Both traffic and R-APS block status). The R-APS channel is never blocked on sub-rings without virtual channel enabled.

Blocked: the status for Port 0 (both traffic and R-APS block status) is ‘blocked’.

Unblocked: the status for Port 0 (both traffic and R-APS block status) is ‘unblocked’.

Port 1 Block Status

Block status for Port 1 (Both traffic and R-APS block status). The R-APS channel is never blocked on sub-rings without virtual channel enabled.

Blocked: the status for Port 1 (both traffic and R-APS block status) is ‘blocked’.

Unblocked: the status for Port 1 (both traffic and R-APS block status) is ‘unblocked’.

FOP Alarm

Displays the Failure of Protocol Defect (FOP) status. If FOP is detected, the red LED (●) displays in the table; the green LED (●) displays if FOP is not detected.

Due to errors in provisioning, the ERP Control Process may detect a combination of conditions which should not occur during "normal" conditions. To warn of such an event, a Failure of Protocol – Provisioning Mismatch (FOP-PM) is defined. The FOP-PM defect, detected if the RPL Owner Node receives one or more No Request R-APS messages with the RPL Blocked status flag set (NR, RB), and a Node ID that differs from its own. The ERP Control Process must notify the equipment fault management process when it detects such a defect condition, and will continue its operation as well as possible. This is only an overview of the defect condition. The associated defect and its details are defined in ITU-T G.8021 as amended by ITU-T G.8021 Amd.1 and Amd.2. Other than alarm noting the defect condition, the ERP state machine continues operation as well as possible.

Buttons

Save: Click to save changes.

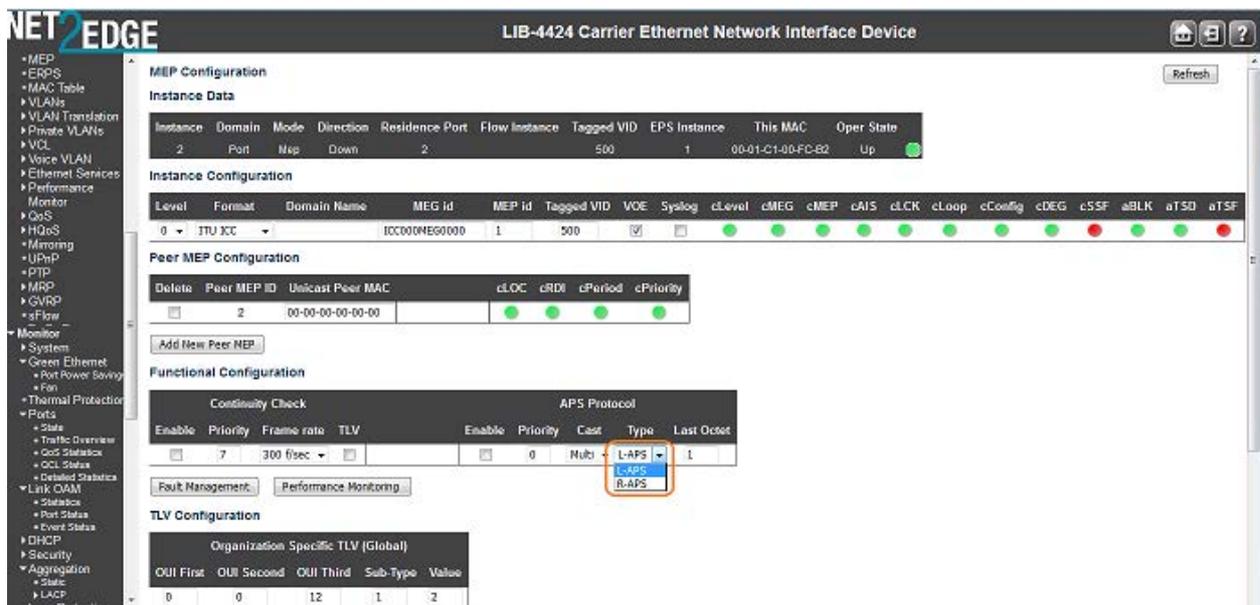
Reset: Click to undo any changes made locally and revert to previously saved values.

Refresh: Click to refresh the page immediately.

Auto-refresh: Click to automatically refresh the page every three seconds.

Ring Protection and MEP Configuration

The LIB-44xx lets you configure the RPL port so it can act in the role of Owner or Neighbour on that ring instance. The WTR time and Hold off timer serve the same purpose as EPS. The Guard timer on the EPRS instance is configurable and helps in ignoring aged R-APS messages that circulate around the ring. The Ring ports have the Y.1731 MEPs which exchange CCMs to monitor the health of the link and also trigger signal failures that can cause the link failure and protection to activate.



The screen above shows a MEP configured with R-APS as the APS Protocol Type. See the “[MEP Configuration](#)” section on page 156 for more information.

Ring Protection Conditions and Commands

The LIB-44xx supports the Ethernet ring **SF** and **NR** conditions per Rec. ITU-T G.8032/Y.1344 (03/2010).

Signal fail (SF) - When an SF condition is detected on a ring link, and it is determined to be a "stable" failure, Ethernet ring nodes adjacent to the failed ring link initiate the protection switching mechanism described in the ITU Recommendation.

No request (NR) - The condition when no local protection switching requests are active.

The **FS**, **MS**, and **Clear** administrative commands are supported:

Forced switch (FS) - This command forces a block on the ring port where the command is issued.

Manual switch (MS) - In the absence of a failure or FS, this command forces a block on the ring port where the command is issued.

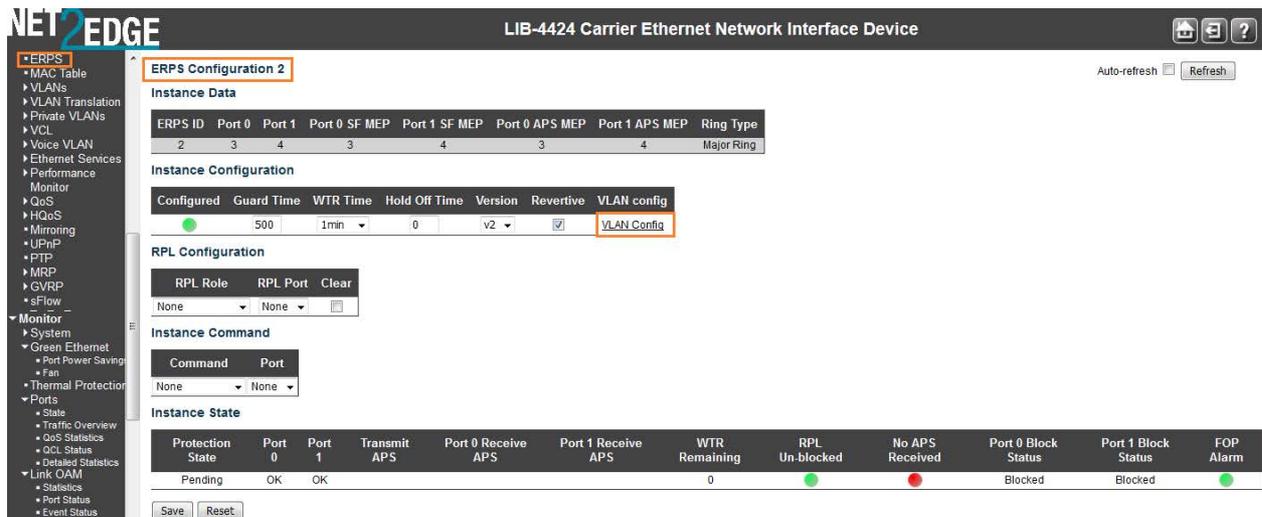
Clear - The Clear command is used for these operations:

- a) Clearing an active local administrative command (e.g., Forced switch or Manual switch).
- b) Triggering reversion before the WTR or WTB timer expires in case of revertive operation.
- c) Triggering reversion in case of non-revertive operation.

Note that at the time of this publication, other commands (Lockout of protection, Replace the RPL, Exercise signal) are undergoing further ITU-T study.

ERPS VLAN Configuration

VLAN config: click the [VLAN Config](#) link in the ERPS Instance Configuration table to display the related ERPS VLAN Configuration page.



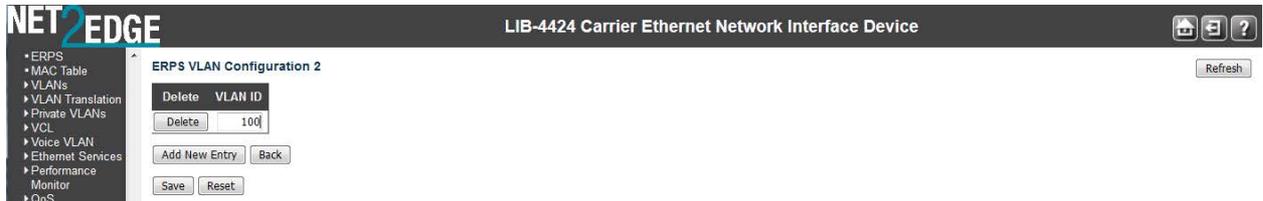
ERPS VLANs are used because Ethernet ring protection configured as a single instance only works at the physical level (adjacent nodes must be directly connected). The ring protection operates at the interface (port) level and not at the VLAN level.

The VLANs created here are tied to Ring instances that are like traffic channels that contain different sets of VLANs. A ring instance is responsible for the protection of a subset of VLANs that transport traffic over the physical ring.

When ring instances are configured for the ring, each ring instance should have its own RPL owner, an east and a west interface, and a ring protection link end.



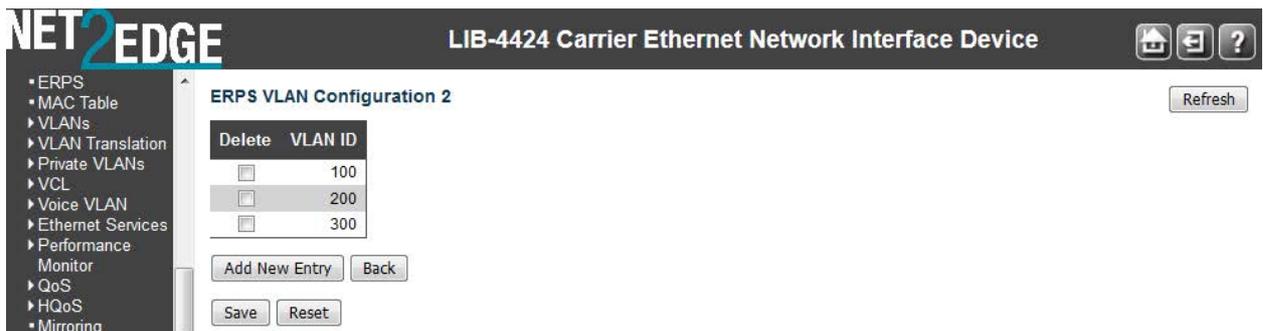
Click the **Add New Entry** button to display the ERPS VLAN Configuration page.



Enter a new ERPS VLAN ID. The default is **0**. The valid range is **1-4094**. Click the **Save** button when done.

Example

Add one or more new VLAN IDs as required.



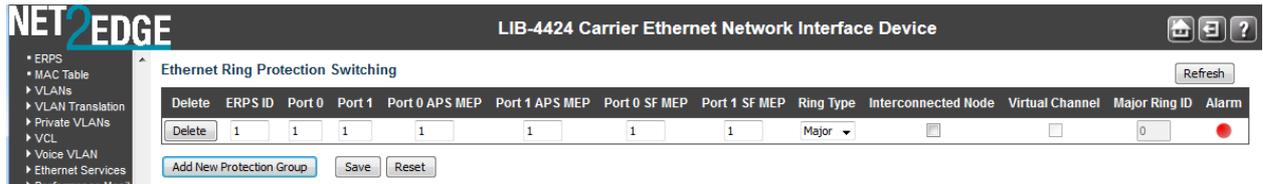
The example above shows ERPS VLAN Configuration 3 with VLAN IDs 1, 2 and 3 configured.

Add a New ERPS Protection Group Procedure

Navigate to the **Configuration > ERPS** menu path. The ERPS table displays.



Click the **Add New Protection Group** button. Entry fields display for the new Protection Group.



Enter unique parameters for ERPS ID, Port 0, Port 1, Port 0 APS MEP, Port 1 APS MEP, Port 0 SF MEP, and Port 1 SF MEP.

At the **Ring Type** dropdown, select **Major** or **Sub**.

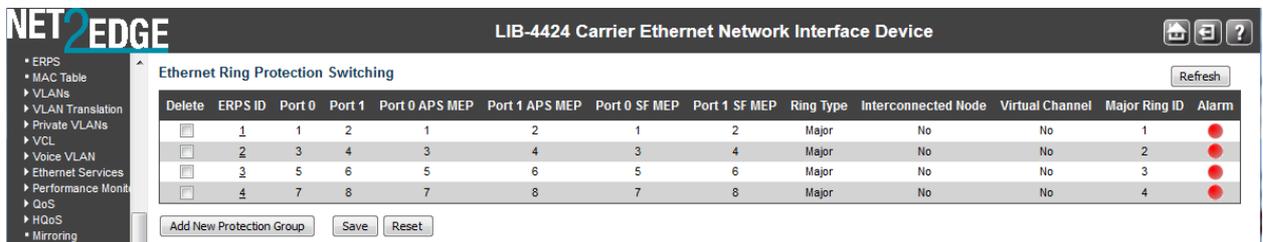
Check or uncheck the **Interconnected Node** and **Virtual Channel** checkboxes.

Enter a **Major Ring ID** (only applies if **Ring Type** = **Sub** was selected in step 4 and Interconnected Node was checked in step 5).

When done click the **Save** button.

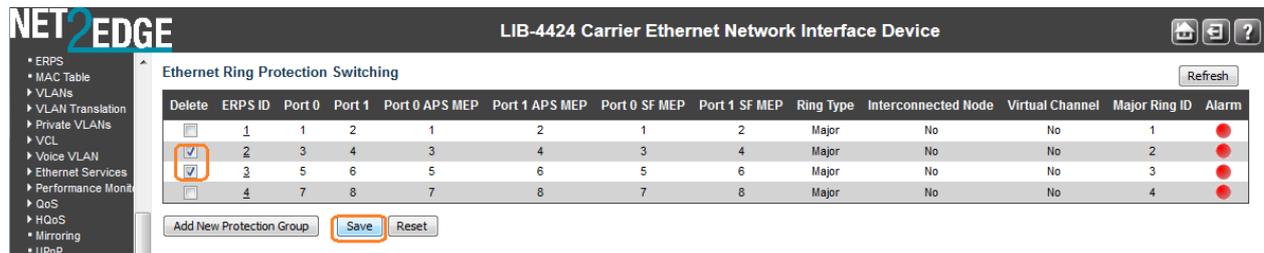
Repeat steps 2-7 for each new Protection Group to be added.

Verify your ERPS configuration (e.g., ERPS IDs 1, 2, 3, and 4 shown below):

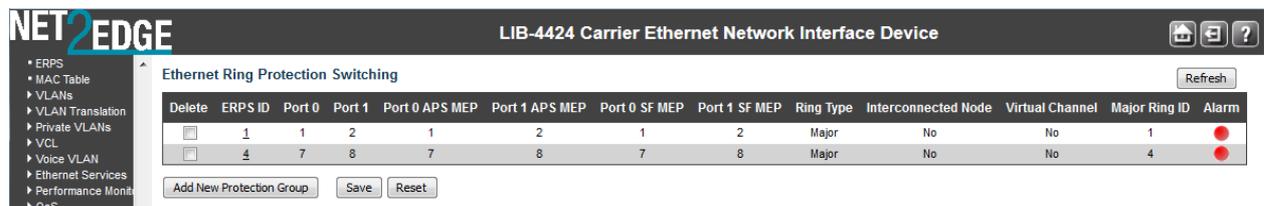


Delete an Existing ERPS Protection Group Procedure

To delete an existing ERPS, check its checkbox in the **Delete** column.



Click the **Save** button.



Verify that the selected ERPS instance(s) are deleted from the table. Click the **Refresh** button if necessary.

MAC Address Table Configuration

The LIB-44xx **Configuration > MAC Table** menu path supports MAC Address table configuration in terms of Aging Configuration, MAC Table Learning, and Static MAC table Configuration.

Switching of frames is based on the DMAC address contained in the frame. The LIB-44xx builds up a table that maps MAC addresses to LIB-44xx ports for knowing to which ports the frames should go, based on the DMAC (Destination MAC) 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 LIB-44xx ports.

The frames also contain a MAC address (SMAC address), which shows the MAC address of the equipment sending the frame. The SMAC (Source MAC) address is used by the LIB-44xx 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 the configured Aging Time.

Enabling known MAC address traffic involves port security. Port security can be either 'static' or 'dynamic'.

Static port security lets you specify which devices are allowed access through a given port. This is done manually by entering the "allowed" device MAC addresses in the MAC address table. Static port security is also known as "MAC address filtering".

Dynamic port security is similar, but instead of specifying the MAC address of the devices, you specify the maximum number of devices to be allowed on the port. If the maximum number that you specify is more than the number of MAC addresses specified manually, the switch learns the MAC

address automatically, up to the maximum specified. If the maximum number specified is less than the number of MAC addresses already specified statically, an error message displays.

The screenshot shows the 'MAC Address Table Configuration' page. On the left is a navigation menu with categories like MAC Table, VLANs, Ethernet Services, and Monitor. The main content area is divided into three sections:

- Aging Configuration:** Contains a checkbox for 'Disable Automatic Aging' (unchecked) and a text input for 'Aging Time' set to '300 seconds'.
- MAC Table Learning:** A table with 29 columns for 'Port Members' (1-29) and three rows for 'Auto', 'Disable', and 'Secure' learning modes. Each cell contains a radio button.
- Static MAC Table Configuration:** A table with columns for 'Delete', 'VLAN ID', 'MAC Address', and 'Port Members' (1-29). Below the table are buttons for 'Add New Static Entry', 'Save', and 'Reset'.

The MAC Address Table is configured on this page. Set timeouts for entries in the dynamic MAC Table and configure the static MAC table here.

Aging Configuration

By default, dynamic entries are removed from the MAC table after 300 seconds (5 minutes). This removal is called 'aging'.

Disable Automatic Aging

Check the checkbox to disable automatic aging.

Aging Time

The FDB is configured with aging time for dynamic learned entries in 1 second increments with the lowest being **10** seconds to a maximum of **1000000** seconds (11.57 days). The default value is **5** minutes (300 seconds). Aging can be disabled by setting the aging time to **0**. Note that when aging is disabled, the FDB size can grow to a maximum of 32768 entries. After the maximum limit of the MAC limit is reached, the new MAC entries are added and the older dynamic MAC entries are purged from the database.

Configure aging time by entering a value here in seconds; for example, **Aging Time 300 seconds** in the screen sample above. The valid values are **0** and **10 - 1000000** seconds (11.57 days).

Disable the automatic aging of dynamic entries by checking the **Disable Automatic Aging** checkbox.

MAC Table Learning

All LIB-44xx learning and switching is based on the MAC forwarding and filtering database (FDB). The LIB-44xx web interface provides a way to purge all dynamic-only or static and dynamic entries out of the FDB. The LIB-44xx FDB can be configured with static MAC entries for filtering or forwarding. The static entries are not aged out and remain in the device even after a power cycle.

The Management interface provides options to add, edit, and delete static entries. Up to 1000 static entries can be stored. Each static entry also has an associated priority which will be used for QoS.

See the 'QoS Configuration' section on page 248 for details on user priority setting.

Bridge ports can be configured to be enabled or disabled for MAC forwarding. When the port is in disabled state, no learning/forwarding takes place. MAC Table Learning can be set to Auto, Disable, or Secure. This information is only available on this page. The MAC Table Learning setting is stored and will be restored after a power cycle.

Per IETF RFC 2233 section 3.1.13, if a port is administratively down, the operational state of the port is also brought down and is not a fault condition. If the administrative state is up but the operational state is down, it implies a fault, and a notification must be sent.

If the learning mode for a given port is greyed out, another module has control of the mode, so that it cannot be changed by a user. An example of such a module is the MAC-based Authentication under 802.1X.

LIB-44xx ports can do learning based on these settings:

Auto

If selected, learning is done automatically as soon as a frame with an unknown SMAC (Source MAC address) is received.

Disable

No learning is done if selected.

Secure

If selected, only static MAC entries are learned, all other frames are dropped.

Note: Make sure that the link used for managing the LIB-44xx 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 LIB-44xx via the serial interface. Static MAC Table Configuration.

The static entries in the MAC table are shown here. The static MAC table can contain up to 64 entries.

The MAC table is sorted first by VLAN ID and then by MAC address.

Delete

Check to delete the entry. It will be deleted during the next Save.

VLAN ID

The VLAN ID of the entry. Each new line (VLAN entry) must have a unique VLAN ID.

MAC Address

Displays the assigned MAC address of the entry.

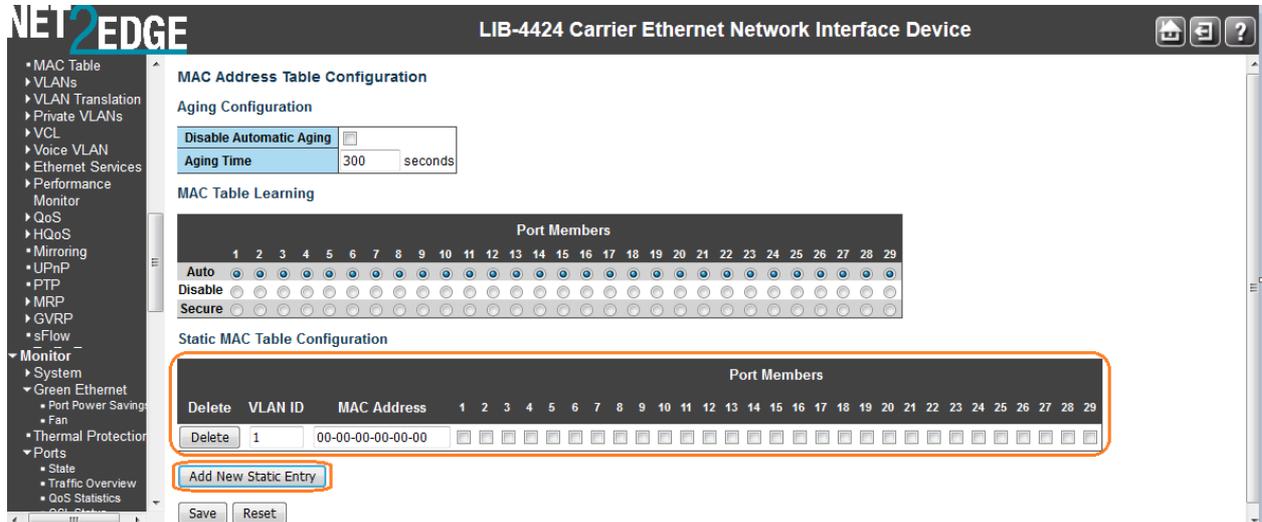
Port Members

Check one or more checkboxes to indicate which port(s) are members of the entry. Check or uncheck as needed to modify the entry. **Note:** If none are checked, the MAC addresses for all ports

will be blocked. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

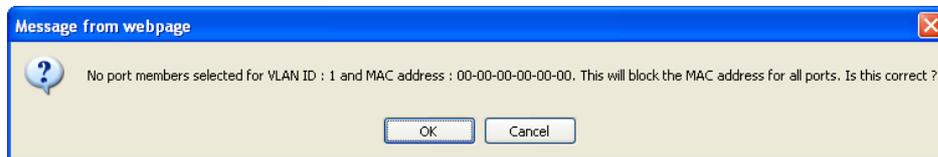
Add New Static Entry

Click the **Add New Static Entry** button to add a new entry to the static MAC table.



Specify the VLAN ID (1-4094), MAC address, and Port Members included for the new MAC entry.

Note: If you do not select any Port Members and then click the **"Save"** button, a message displays warning you that this will block the MAC address for all ports.



If this is the configuration you want, click the **OK** button. Otherwise click the **Cancel** button and select one or more Port Members.

Click the **"Save"** button when done.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

VLAN Translation Configuration

The LIB-44xx lets you configure VLAN translation (mapping) from the **Configuration > VLAN Translation Configuration** menu path. Here you can configure the **Port to Group Mapping** and the **VID Translation Mapping** functions.

Port to Group Mapping

Click the **"Add New Entry"** button to display the entry line of the Port to Group mapping Table.

The screenshot shows the 'VLAN Translation Port Configuration' page. On the left is a navigation menu with options like 'VLAN Translation', 'Private VLANs', 'VCL', 'Voice VLAN', 'Ethernet Services', 'Performance Monitor', 'QoS', 'HQoS', 'Mirroring', 'UPnP', and 'PTP'. The main content area features a table titled 'Group Configuration' with the following structure:

Port	Default	Group ID
*	<input type="checkbox"/>	<>
1	<input type="checkbox"/>	1
2	<input type="checkbox"/>	2
3	<input type="checkbox"/>	3
4	<input type="checkbox"/>	4
5	<input type="checkbox"/>	5
6	<input type="checkbox"/>	6

At the top right of the configuration area, there are 'Auto-refresh' and 'Refresh' buttons.

This page lets you map a set of up to **four** Port members to a Group ID for all LIB-44xx ports. The displayed settings are explained below:

Port

The Port column shows the list of ports 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

A valid Group ID is an integer value from **1** to **4**. A set of VLAN Translations are mapped to a Group ID. This way a port is mapped to a list of VLAN Translations easily by mapping it to a Group. The number of Groups in this LIB-44xx is equal to the number of ports (8) present in this LIB-44xx. A port can be mapped to any of the Groups. Multiple ports can also be mapped to a Group with same Group ID.

Note: By default, each port is mapped to a Group with a Group ID equal to the port number. For example, port 1 is mapped to the Group with ID=1.

Adding a New Port to Group Mapping Entry

Navigate to the **Configuration > VLAN Translation Configuration** menu path.

This is a duplicate of the screenshot above, showing the 'VLAN Translation Port Configuration' page with the same table and navigation elements.

Click a corresponding radio button to make port to be member of a different Group (e.g., Group ID 1).

Click the **Save** button to confirm the changes.

Click the corresponding radio button to make the port a member of the requisite Group.

Port	Default	Group ID
*		<>
1		10
2		2
3		3
4		4
5		5
6		6

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Refresh: Click to refresh the page immediately.

Auto-refresh: Click to automatically refresh the page every three seconds.

VID Translation Mapping

LIB-44xx VID Translation Mapping is done from the **Configuration > VLAN Translation > VID Translation Mapping** menu path. This page lets you map a VLAN ID to other VLAN IDs for a particular Group ID globally.

The VLAN Translation Table initially displays “No VLAN Translation entry found”. Click the “**Add New Entry**” button to display the entry fields.

The displayed settings are explained below:

Delete

To delete a VLAN Translation Group database entry, click the **Delete** button. The entry will be deleted from the LIB-44xx during the next Save.

Group ID

A valid Group ID is an integer value from 1 to 10. A set of VLAN Translations are mapped to a Group ID. This way a port is mapped to a list of VLAN Translations easily by mapping it to a group. The number of groups in a LIB-44xx is equal to the number of ports present in this LIB-44xx. A port can be mapped to any of the groups. Multiple ports can also be mapped to a group with same Group ID.

Note: By default, each port is mapped to a group with a Group ID equal to the port number. For example, port 1 is mapped to the group with Group ID=1.

VLAN ID

Sets / indicates the ID to which Group ID will be mapped. A valid VLAN ID ranges from 1-4094. The VLAN ID is a 12-bit field specifying the VLAN to which the frame belongs. The **VLAN ID** and the **Translated to VID** cannot be the same.

Translated to VID

Sets / indicates the VID to which VLAN ID of ingress frames will be changed, if the VID in incoming frames is the same as configured in VLAN ID field (preceded by this field) on member ports of a particular group to which this entry belongs. The **VLAN ID** and the **Translated to VID** cannot be the same.

Add New Entry (to the VLAN Translation table)

Click the 'Add New Entry' button to add a new entry in the VLAN Translation table. An empty row is added to the table, the Group ID, VLAN ID and Translated to VID fields can be configured as needed. Valid values for a VLAN ID are **1** through **4094**.

You can use the **Delete** button to undo the addition of new entry.

Buttons

Auto-refresh: Click to automatically refresh the page every three seconds.

Refresh: Click to refresh this page immediately.

Add New Entry: Click to add a new entry in VLAN Translation table. See above.

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

When the Save is complete, the new VLAN Translation entry displays in the table.

The screenshot shows the 'VLAN Translation Mapping Table' configuration page. The table has the following data:

Group ID	VID	TVID	
10	30	10	⊕ ⊗
20	40	20	⊕ ⊗

VLANs Configuration

The LIB-44xx lets you configure VLANs from the **Configuration > VLANs** menu path. Here you can configure the 'VLAN Membership' and the 'Ports' sub-menu functions.

A VLAN (Virtual LAN) is a method to restrict communication between switch ports. VLANs can be used for the following applications:

VLAN unaware switching: This is the default configuration. All ports are VLAN unaware with Port VLAN ID 1 and members of VLAN 1. This means that MAC addresses are learned in VLAN 1, and the switch does not remove or insert VLAN tags.

VLAN aware switching: This is based on the IEEE 802.1Q standard. All ports are VLAN aware. Ports connected to VLAN aware switches are members of multiple VLANs and transmit tagged frames. Other ports are members of one VLAN, set up with this Port VLAN ID, and transmit untagged frames.

Provider switching: This is also known as Q-in-Q switching. Ports connected to subscribers are VLAN unaware, members of one VLAN, and set up with this unique Port VLAN ID. Ports connected to the service provider are VLAN aware, members of multiple VLANs, and set up to tag all frames. Untagged frames received on a subscriber port are forwarded to the provider port with a single VLAN tag.

Tagged frames received on a subscriber port are forwarded to the provider port with a double VLAN tag.

The LIB-44xx is compliant with IEEE 802.1Q standard. The LIB-44xx is capable of VLAN bridging and filtering. By default the devices comes up with all ports belonging to the same VLAN.

The LIB-44xx supports independent VLAN learning (IVL) by maintaining different FDB for each VLAN as per IEEE802.1Q Appendix A. This helps to maintain separate FDB for each VLAN. The usefulness of IVL is explained in IEEE802.1Q Appendix A.

Each VLAN is identified by a VLAN ID, a 12-bit field specifying the VLAN to which the frame belongs. The LIB-44xx supports the entire range of 4k VLAN IDs except for the following:

- a. VLAN ID = 0 is used for priority tagged traffic hence will not be used.
- b. VLAN ID = 1 is used for default VLAN on the device (native VLAN ID)
- c. VLAN ID = 4095 is reserved and not available for normal traffic.

Each VLAN has a unique string for identification called the “VLAN Name”, no spaces will be allowed. A maximum of 64 VLANs can have VLAN names and each name is restricted to 32 bytes. Only alpha and numeric digits are allowed as valid characters with at least one alpha character is required as part of the name.

VLANs > Ports

LIB-44xx VLAN port configuration is done from the **Configuration > VLANs > Ports** menu path. This page is used for configuring the LIB-44xx port VLAN. Also, The Management Port is configured at "**Configuration > VLANs > Ports: Management Port - PortType**".

The LIB-44xx ports can be configured with a default or native VLAN ID, so that all untagged and priority tagged traffic will be classified to this VLAN ID. The native VLAN ID for all ports is set to 1 but is user configurable on a per port basis. Hence all ports by default belong to the same broadcast domain.

The screenshot shows the NET2EDGE web interface for a LIB-4424 Carrier Ethernet Network Interface Device. The 'Global VLAN Configuration' section includes 'Allowed Access VLANs' set to 1 and 'Ethertype for Custom S-ports' set to 88A8. The 'Port VLAN Configuration' table is as follows:

Port	Mode	Port VLAN	Port Type	Ingress Filtering	Ingress Acceptance	Egress Tagging	Allowed VLANs	Forbidden VLANs
1	Access	1	C-Port	<input checked="" type="checkbox"/>	Tagged and Untagged	Untag All	1	
2	Access	1	C-Port	<input checked="" type="checkbox"/>	Tagged and Untagged	Untag All	1	
3	Access	1	C-Port	<input checked="" type="checkbox"/>	Tagged and Untagged	Untag All	1	
4	Access	1	C-Port	<input checked="" type="checkbox"/>	Tagged and Untagged	Untag All	1	
5	Access	1	C-Port	<input checked="" type="checkbox"/>	Tagged and Untagged	Untag All	1	
6	Access	1	C-Port	<input checked="" type="checkbox"/>	Tagged and Untagged	Untag All	1	

The VLAN Port Configuration table parameters are explained below:

Ethertype for Custom S-ports 0x

This field specifies the Ether type used for Custom S-ports. This is a global setting for all of the Custom S-ports. The valid ‘EtherType’ range is restricted to 0x600 to FFFF. Typical ‘EtherType’ values include:

8100: VLAN-tagged frame (IEEE 802.1Q).

88A8: Provider Bridging (IEEE 802.1ad).

9100: VLAN Tag Protocol Identifier (Q-in-Q).

Using Provider mode (Provider EtherTypes X8100, X9100, X88A8) will ensure that the frame egressing this port will always have the S-Tag added, regardless of whether the ingress frame was already C-Tagged or plain untagged.

When the 0x8100 tag is added twice, the outer tag is called the Provider tag and the inner one is called the Customer IEEE 802.1Q tag. The inner VLAN tag is referred to as the customer VLAN tag (C-Tag) because the customer assigns it. Before the standardization, some vendors used 0x8100 and 0x9100 for outer Provider tagging. The 0x88A8 tag was adapted by the IEEE later.

Select **X8100**, or **X9100**, or **X88A8**. In Provider mode, a frame is considered Provider tagged if it matches the 'Provider Ether Type'. Frames that are ingress with a Provider tag are stripped of their Provider tag on egressing this interface. If the frame's Ethertype doesn't match the Provider Ether Type it is considered as untagged. The default is **x88a8**. See [Glossary](#) entry "Commonly Used EtherTypes" for more information.

Management Port - PortType

The LIB-44xx port type selection for the Management Port; either **Unaware**, **C-port**, **S-port**, or **S-custom-port** where:

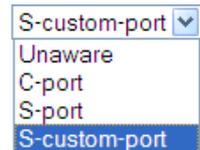
Unaware: all frames are classified to the Port VLAN ID and tags are not removed (the default setting).

C-port: Customer port; •with Customer tags (C-Tags), which use TPID 0x8100.

S-port: Service port; with Service tags (S-Tags), which use TPID 0x88A8 (IEEE 802.1ad).

S-custom-port: Custom Service port; with Service tags (S-Tags), which use a custom TPID assigned in the LIB-44xx firmware.

For Customer tags and Service tags, both VLAN tags (tags with non-zero VID) and Priority tags (tags with VID = 0) are processed.



The tag header is either retrieved from a tag in the incoming frame or from a default port-based tag header. The port-based tag header is configured in the LIB-44xx firmware.

For double-tagged frames, there is an option to use the inner tag instead of the outer tag.

In addition to the tag header, the ingress port decides the number of VLAN tags to pop at egress. If the configured number of tags to pop is greater than the actual number of tags in the frame, the number is reduced to the actual number of tags in the frame.

Port

This is the logical port number of this row (e.g., port s1-4) on the LIB-4400. The * in the Port column acts as a 'wild card' character which causes the selections in this row to be applied to all other Ports (rows) in the table for which this selection is valid. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Port Type

A Port can be one of the following types: **Unaware**, **C-port**, **S-port**, or **S-custom-port** where:

Unaware: all frames are classified to the Port VLAN ID and tags are not removed (the default setting).

C-port: Customer port; •with Customer tags (C-Tags), which use TPID 0x8100.

S-port: Service port; with Service tags (S-Tags), which use TPID 0x88A8 (IEEE 802.1ad).

S-custom-port: Custom Service port; with Service tags (S-Tags), which use a custom TPID assigned in the LIB-44xx firmware.

For Customer tags and Service tags, both VLAN tags (tags with non-zero VID) and Priority tags (tags with VID = 0) are processed.

The tag header is either retrieved from a tag in the incoming frame or from a default port-based tag header. The port-based tag header is configured in the LIB-44xx firmware.

For double-tagged frames, there is an option to use the inner tag instead of the outer tag.

In addition to the tag header, the ingress port decides the number of VLAN tags to pop at egress. If the configured number of tags to pop is greater than the actual number of tags in the frame, the number is reduced to the actual number of tags in the frame.

Ingress Filtering

Enable ingress filtering on a port by checking this checkbox. This parameter affects VLAN ingress processing.

If ingress filtering is enabled and the ingress port is not a member of the classified VLAN of the frame, the frame is discarded. By default, ingress filtering is disabled (no checkmark).

Frame Type

Defines whether the port accepts all frames, or only tagged frames, or only untagged frames.

This parameter defines VLAN ingress processing:

If **Tagged** is selected, the port only accepts tagged frames, and untagged frames received on the port are discarded.

If **Untagged** is selected, the port only accepts untagged frames, and tagged frames received on the port are discarded.

If **All** is selected, both tagged and untagged frames are accepted on this port and no received tags are discarded. By default, the Frame Type field is set to **All**.

Port VLAN Mode

Configures the Port VLAN Mode. The allowed values are **None** or **Specific**. This parameter affects VLAN ingress and egress processing.

If **None** is selected, a VLAN tag with the classified VLAN ID is inserted in frames transmitted on the port. This mode is normally used for ports connected to VLAN aware switches.

If **Specific** (the default value) is selected, a Port VLAN ID can be configured (see below). Untagged frames received on the port are classified to the Port VLAN ID. If VLAN awareness is disabled, all frames received on the port are classified to the Port VLAN ID. If the classified VLAN ID of a frame transmitted on the port is different from the Port VLAN ID, a VLAN tag with the classified VLAN ID is inserted in the frame.

Port VLAN ID

Configures the port VLAN identifier (PVID). The allowed values are **1** through **4094**. The default value is **1**.

Note: The port must be a member of the same VLAN as the Port VLAN ID. If the Port VLAN Mode is set to "None", then the Port VLAN ID field is greyed out and no entry is allowed.

Note: configure the Management VLAN before configuring the Management Port here. See "[Configuration > System > IP > VLAN ID](#)" on page 12.

The Management Port is configured at "[Configuration > VLANs > Ports: Management Port - PortType](#)".

Tx Tag

Determines egress tagging of a port, where:

Untag_pvid - All VLANs except the configured PVID will be tagged.

Tag_all - All VLANs are tagged.

Untag_all - All VLANs are untagged.

Buttons

Auto-refresh: Check this box to refresh the page automatically every 3 seconds.

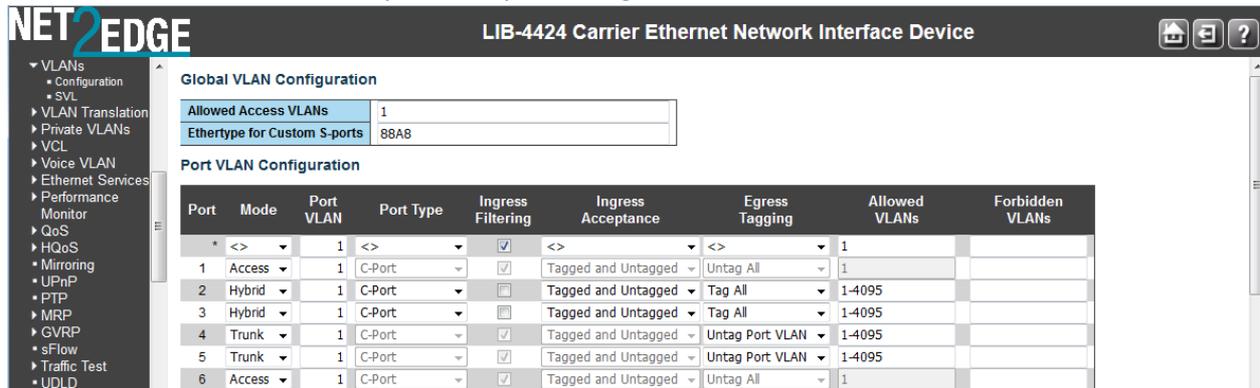
Refresh: Click to refresh the page immediately.

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

When done, review your VLAN Port configuration. Click the **Save** button when finished.

The screen below shows a sample VLAN port configuration.



VID (VLAN ID) Range Summary

Some LIB-44xx modules accept VID range 0 – 4094 and some modules only accept VID range 1 – 4094.

A VID of 0 is only for modules where untagged/priority tag makes sense. LLDP and LOAM are untagged. The valid VID ranges are summarized below:

Module	Valid VID Range
ACL	(1-4094)
ARP Inspection	(1-4094)
ERPS	(1-4094)
EVC	
1. EVC parameters: 1) VID (1-4094), 2) IVID(1-4094)	
Inner Tag-->VLAN ID (0-4094)	
Outer Tag-->VLAN ID (0-4094)	
2. ECEs > UNI Matching > VLAN ID Value (0-4095)	
IP (Management VLAN)	(1 -4094)
IP Source Guard	(1 - 4094)
IPMC	(1-4094) (only up to VLANs)
LLDP	(none)
MAC	(1-4094)
MEP UP	(1-4094) (depends on the VLANs)
MEP Down	(0-4094)
MVR	(1-4094)
PTP	(0-4094)
QOS (QCL)	(1-4094)
VCL	(1-4094)

VLAN translation (1-4094)

Note: when you edit the "Inner Tag VLAN ID Range" and you select the "Range" in the Inner VLAN ID Filter, if the expected maximum number of VIDs is 4094, you cannot apply the ECE entry with a valid range VLAN of 2048. The default VID range is 0-2047 in this case.

Provider Bridging (IEEE 802.1ad 2005)

The LIB-44xx is compliant with IEEE802.1ad 2005 standard in recognizing S-Tags and C-Tags. The LIB-44xx can recognize S-Tag or C-Tag based on Ethertype. The default value is 0x88a8 for S-Tag and 0x8100 for C-Tag as per the standard. You can configure the Ethertype of S-Tag via the web interface. A default list of 0x88a8, 0x8100 and 0x9100 is provided on the web interface.

Tag Type	Name	Value
Customer VLAN tag	IEEE 802.1Q Tag Protocol Type (802.1Q Tag Type)	8100
Service VLAN tag	IEEE 802.1Q Service Tag Type (802.1Q S-Tag Type)	88A8
Q-in-Q	VLAN Tag Protocol Identifier	9100

Figure 11. 802.1Q EtherTypes (excerpt from IEEE 802.1ad 2005)

The LIB-44xx can be an SVLAN bridge, C-bridge, or both bridge types, and the hardware can support inspection of both the tags. The LIB-44xx can push and pop one or both tag types. Provider Tagging Use cases

All to one bundling services: In this scenario the device caters to multi-tenant services. This service is the Ethernet private line as stated by the MEF 30 standard.

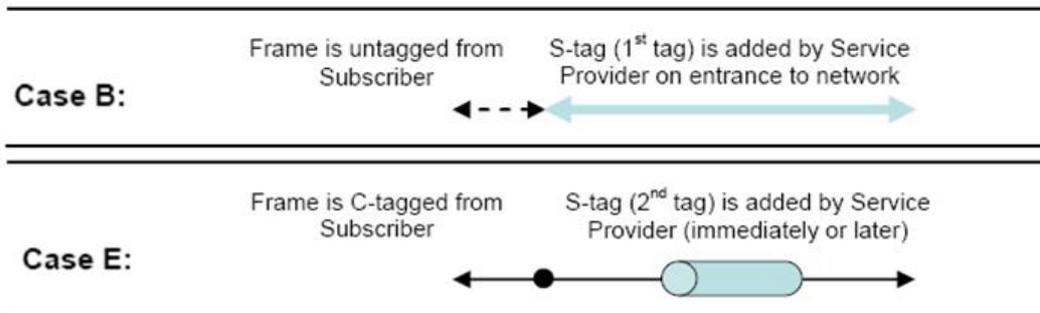


Figure 12. All to one bundling VLAN Cases

All traffic from customer facing ports is bundled using an S-VLAN tag. The SVID will identify the customer traffic within the provider network; at the hand-off on the other side of the network, the S-tag is stripped. The customer traffic can be Untagged, Priority tagged or C-VLAN tagged; in some cases it can be all.

The provider S-tag bundles all transparently based on the ingress port and transported through the provider network, hence the name 'all to one bundling'.

Private VLANs Configuration

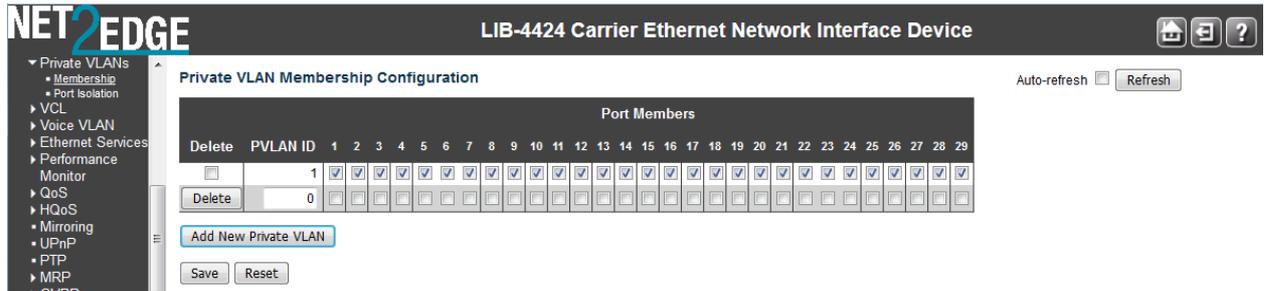
The LIB-44xx lets you configure Private VLANs (PVLANS) from the **Configuration > Private VLANs** menu path. Here you can configure the 'PVLAN Membership' and the 'Port Isolation' sub-menu functions.

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.

PVLAN Membership

A VLAN can be configured as a Private VLAN (PVLAN). In a PVLAN, communication between ports within that private VLAN is not permitted. The LIB-44xx offers more control on forwarding/routing frames using this feature. Private VLANs let you create a forwarding ‘mask’ on a VLAN on which a port can communicate with other ports belonging to a VLAN. A port must be a member of both the 802.1Q VLAN and the private VLAN to forward frames.

From the default page, click the “Add New Private VLAN” button to display the entry table.



This page lets you:

- Monitor and modify LIB-44xx Private VLAN membership configurations,
- Add or delete Private VLANs, and
- Add or delete Private VLAN Port Members.

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.

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.

The VLAN Port Membership configuration table parameters are explained below:

Delete

To delete a private VLAN entry, check this box. The entry will be deleted during the next Save.

PVLAN ID

Indicates the ID of this particular private VLAN. Enter a unique number from 1-4.

Port Members

Displays a row of checkboxes for each port 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 checkboxes are unchecked.

Member ports of a PVLAN can communicate with each other. Isolated ports configured as part of a PVLAN cannot communicate with each other.

Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Add New Private VLAN

Click the **Add New Private VLAN** button to add a new private VLAN ID. An empty row is added to the table, and the private VLAN can be configured as needed. The allowed range for a private VLAN ID is the same as the LIB-44xx port number range. Any values outside this range are not accepted, and a warning message appears. Click "OK" to discard the incorrect entry, or click "Cancel" to return to the editing and make a correction.

The Private VLAN is enabled when you click the "Save" button.

You can use the **Delete** button to undo the addition of new Private VLANs.

Buttons

Auto-refresh: Check this box to refresh the page automatically every 3 seconds.

Refresh: Click to refresh the page immediately.

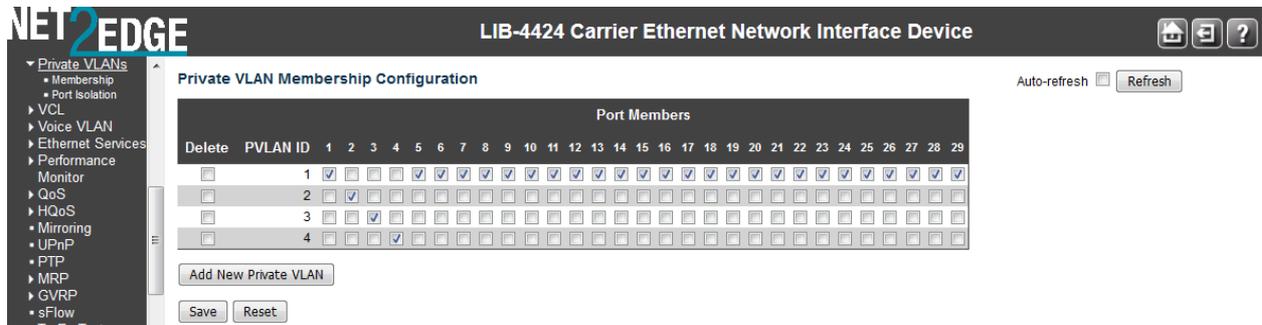
Add New Private VLAN: Click to add a new private VLAN. See above.

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

When done, verify your **Configuration > Private VLANs > PVLAN Membership** configuration.

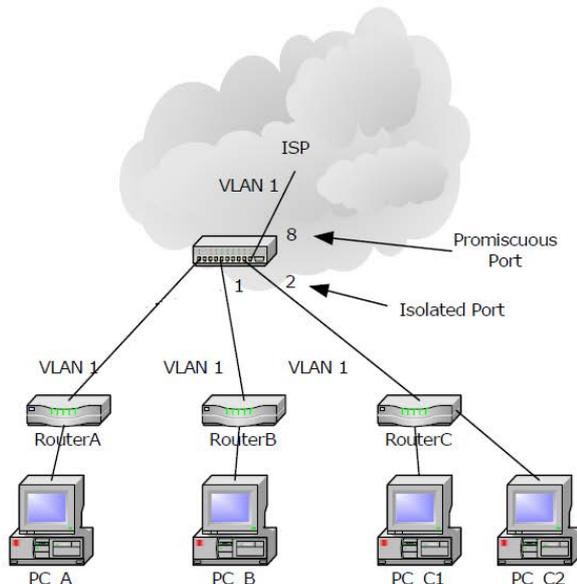
For example, the screen below shows a PVLAN Membership configured with four PVLAN IDs (1-4).



Port Isolation

You can configure port isolation from the **Configuration > Private VLANs > Port Isolation** menu path. This page is used for enabling or disabling port isolation on ports in a Private VLAN.

Port isolation offers isolation of that Port from the VLAN forwarding on the VLAN that it is a member of. Isolated ports configured as part of a PVLAN cannot communicate with each other. Member ports of a PVLAN are not isolated and can communicate with each other.



VLAN Table

VID	Ports in Mask	Private
...
1	0, 1, 2, 8	1
...

ISOLATED_PORTS (Isolated ports set to 0)

Bit	Value
0	1
1	1
2	0
...	...
8	1
...	...

Figure 13. Port Isolation

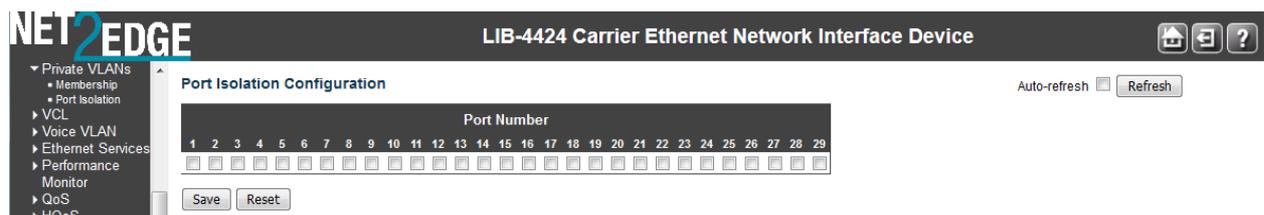
Isolated ports can only receive traffic from promiscuous ports and can send only to promiscuous ports that are part of the PVLAN. ‘Community’ ports are not supported. If ports 2 and 3 are both isolated and members of the same PVLAN then they should not be able to communicate. For Private VLANs to be applied, the switch must be configured for standard VLAN operation. When this is done, one or more of the configured VLANs can be configured as private VLANs. Ports in a Private VLAN can be either **Promiscuous ports** (from which traffic can be forwarded or received from to all ports in the Private VLAN) or **Isolated ports** (ports from which traffic can only be forwarded to or received from promiscuous ports in the Private VLAN).

The configuration of promiscuous and isolated VLANs applies to all private VLANs.

The forwarding of frames classified to a private VLAN happens:

- a) when traffic comes in on a promiscuous port in a private VLAN, the VLAN mask from the VLAN table is applied.
- b) when traffic comes in on an isolated port, the Isolated Port mask is applied in addition to the VLAN mask from the VLAN table.

The default Port Isolation Configuration page is shown below:



A port member of a VLAN can be isolated to other isolated ports on the same VLAN and Private VLAN.

Port Number

A checkbox is provided for each port of a private VLAN (e.g., 1-4 on the LIB-4400).

When checked, port isolation is enabled on that port, and that port is considered “isolated”. Isolated ports configured as part of a PVLAN cannot communicate with each other.

When unchecked, port isolation is disabled on that port, and that port is considered a “member port”. Member ports of a PVLAN are not isolated and can communicate with each other.

By default, port isolation is disabled (unchecked) on all ports (i.e., all ports are “member ports” at default).

Note: At least one port entry must be selected in order to add a new entry. Note that the FPGA port (port 12 or LIB-4424 port 24) is “hidden” when the Shared port is set to Internal mode.

Buttons

Auto-refresh: Check this box to refresh the page automatically every 3 seconds.

Refresh: Click to refresh the page immediately.

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

VCL (VLAN Control List)

The **Configuration > VCL** (VLAN Control List) menu path lets you configure the LIB-44xx for MAC-based VLAN mappings and Protocol-based VLAN mapping. A VCL is used for assigning a particular flow to a particular VLAN. VCLs can enforce VLAN security that is based on a variety of information.

The IND-328x VCL (VLAN Control List) commands let you configure the IND-328x for MAC-based VLAN, Protocol-based VLAN, and/or IP Subnet-based VLAN mappings.

MAC-based VLANs let you add and delete MAC-based VLAN entries and assign the entries to different ports. This page shows only static entries.

Protocol-based VLANs let you configure the LIB-44xx for Protocol-to-Group and/or Group-to-VLAN settings.

IP Subnet-based VLANs let you define a VLAN membership by the subnet to which a device's IP address belongs. You can add, update, and delete IP subnet-based VLAN entries and assign the entries (membership) to different ports. VLANs are layer 2 constructs, compared with IP subnets which are layer 3 constructs. In an environment employing VLANs, a one-to-one relationship often exists between VLANs and IP subnets, but it is possible to have multiple subnets on one VLAN. VLANs and IP subnets provide independent Layer 2 and Layer 3 constructs that map to each other, and this correspondence helps in network design. Note that this involves only static entries.

MAC-based VLAN

The MAC-based VLAN entries can be configured here. This page lets you add and delete MAC-based VLAN entries and assign the entries to different ports. This page shows only static entries. From the default page (“*Currently no entries present*”), click the **Add New Entry** button to display the entry fields.

The screenshot shows the NET2EDGE web interface for a LIB-4424 Carrier Ethernet Network Interface Device. The main configuration area is titled "MAC-based VLAN Membership Configuration". It features a table with columns for "Delete", "MAC Address", "VLAN ID", and "Port Members" (ports 1 through 29). The "Delete" column has a "Delete" button. The "MAC Address" column contains "00-00-00-00-00-00". The "VLAN ID" column contains "1". The "Port Members" column has 29 checkboxes, all of which are currently unchecked. Below the table, there is an "Add New Entry" button, and at the bottom, "Save" and "Reset" buttons. A sidebar on the left contains a navigation menu with options like VCL, MAC-based VLAN, Protocol-based VLAN, IP Subnet-based VLAN, Voice VLAN, Ethernet Services, Performance Monitor, QoS, and Mirroring. The top right of the interface includes an "Auto-refresh" checkbox and a "Refresh" button.

The MAC-based VLAN entries are explained below:

Delete

To delete a MAC-based VLAN entry, check this box and click 'Save'. The selected entry will be deleted.

MAC Address

Enter the MAC address in the format xx-xx-xx-xx-xx-xx.

VLAN ID

Enter the VLAN ID (VID) in the range of 1-4094.

Port Members

A row of check boxes for each port is displayed for each MAC-based VLAN entry. To include a port in a MAC-based VLAN, check the box (). To remove or exclude the port from the MAC-based VLAN, make sure the box is unchecked (). By default, no ports are members, and all boxes are unchecked. At least one port must be checked to add an entry before you click the **'Save'** button. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Add New Entry (MAC-based VLAN Member)

Click the **Add new entry** button to add a new MAC-based VLAN entry. An empty row is added to the table, and the MAC-based VLAN entry can be configured as needed. Any unicast MAC address can be configured for the MAC-based VLAN entry. No broadcast or multicast MAC addresses are allowed. Legal values for a VLAN ID are **1** through **4094**.

The MAC-based VLAN entry is enabled when you click the "Save" button. A MAC-based VLAN without any port members will be deleted when you click the "Save" button.

The **Delete** button can be used to undo the addition of new MAC-based VLANs.

Buttons

Refresh: Refreshes the displayed table.

Auto-refresh: Check this box to refresh the page automatically every 3 seconds.

|<<: Updates the table starting from the first entry in the MAC-based VLAN Table.

>>: Updates the table, starting with the entry after the last entry currently displayed.

Delete: Click to undo the addition of new MAC-based VLANs. The maximum possible MAC-based VLAN entries are limited to 256.

Add New Entry: Click to add a new MAC-based VLAN entry. See above.

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Messages

Message: *MAC address to VLAN ID mapping already exists and it has to be deleted if new mapping for MAC address to VID is required.*

Meaning: Too many mappings exist.

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Delete one or more mappings, Save, and continue operation.

Message: No multicast or broadcast address allowed.

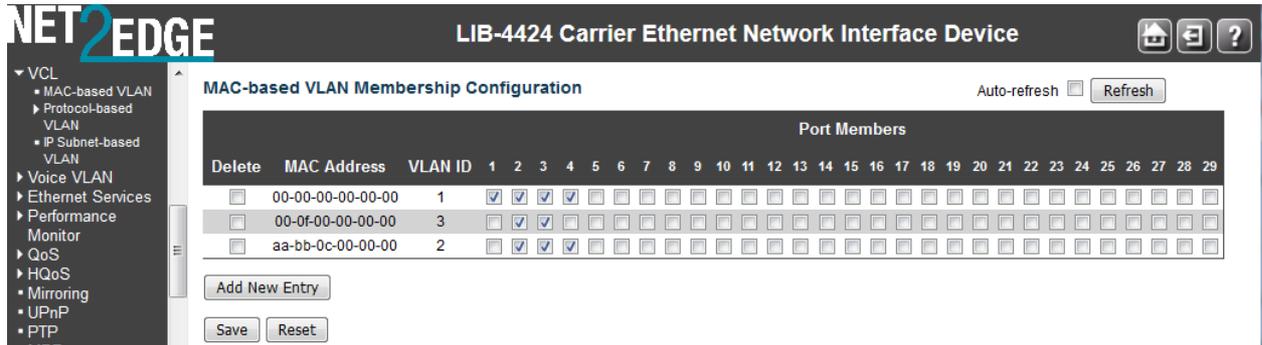
Meaning: You entered an invalid value in the MAC address field (e.g., 77-00-00-00-00-00).

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Enter a valid MAC address and continue operation.

Example

The example below shows three MAC-based VLAN Membership configurations.



In this example above:

VLAN ID 100 has MAC address 00-00-00-00-00-00 with all four ports (1-4) as Port Members.

VLAN ID 200 has MAC address aa-88-09-0d-0f-ee with ports 2 - 4 as Port Members.

VLAN ID 300 has MAC address aa-88-09-0d-0f-ee with ports 2 and 3 as Port Members.

Protocol-based VLAN

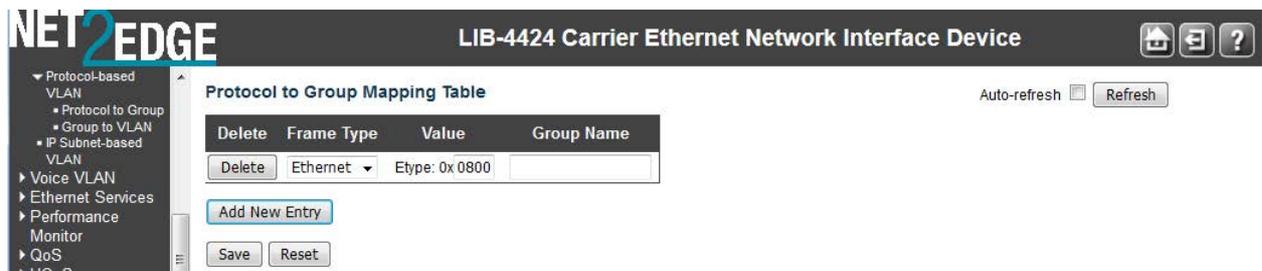
The **Configuration > VCL > Protocol-based VLAN** menu path lets you configure the LIB-44xx for Protocol to Group and/or Group to VLAN settings.

Protocol to Group

This page lets you add new protocols to Group Name (unique for each Group) mapping entries and lets you to view and delete already mapped entries for the LIB-44xx.

At default, the table displays “No Group entry found!”.

Click the **Add New Entry** button to display the entry fields.



The displayed settings are explained below:

Delete

To delete a Protocol to Group Name map entry, check this checkbox. The entry will be deleted on the LIB-44xx at the next ‘Save’.

Frame Type

The different frame types have different formats and MTU values, but can coexist on the same physical medium. Select a ‘Frame Type’ of one of the following values from the dropdown:

Ethernet : the frame type is Ethernet. An Etype value and a Group Name entry are required. The Ethernet frame (the most common type, used directly by the IP).

SNAP : Subnetwork Access Protocol (SNAP) frame. Ethernet SNAP is similar to 802.2 with LLC parameters, but with expanded LLC capabilities. Ethernet SNAP can support IPX/SPX, TCP/IP, and AppleTalk Phase 2 protocols.

LLC : IEEE 802.2 Logical Link Control (LLC) frame. LLC addressing involves LLC protocol data units (PDUs) which contain addressing information, consisting of two fields; the Destination Service Access Point (DSAP) address field, and the Source Service Access Point (SSAP) address field. Each of these is an 8-bit field made up of two components.

Note: On changing the Frame type field, the valid values of the following text field will vary depending on the new frame type you select here.

Value

The valid value that you can enter in this text field depends on the option selected from the preceding 'Frame Type' selection menu. The criteria for the three Frame Types are explained below:

Ethernet: Values in the text field when Ethernet is selected as a Frame Type is called 'Etype'. Valid values for Etype ranges from 0x0600 - 0xffff. The default is Etype: 0x0800.

Protocol to Group Mapping Table

Delete	Frame Type	Value	Group Name
Delete	Ethernet	Etype: 0x0800	

SNAP: Valid value in this case also is comprised of two different sub-values.

- a. **OUI:** OUI (Organizationally Unique Identifier) in the format of xx-xx-xx where each pair (xx) in string is a hexadecimal value ranges from 0x00-0xff.
- b. **PID:** If the OUI is hexadecimal 000000, the protocol ID is the Ethernet type (EtherType) field value for the protocol running on top of SNAP; if the 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 value of OUI field is 00-00-00 then value of PID will be etype (0x0600-0xffff) and if value of OUI is other than 00-00-00 then valid value of PID will be any value from 0x0000 to 0xffff.

Protocol to Group Mapping Table

Delete	Frame Type	Value	Group Name
Delete	SNAP	OUI: 0x00-E0-2B PID: 0x0001	

LLC: Valid value in this case is comprised of two different sub-values.

- a. **DSAP:** 1-byte long string (0x00-0xff). The default is DSAP: 0xFF.
- b. **SSAP:** 1-byte long string (0x00-0xff). The default is SSAP: 0xFF.

Protocol to Group Mapping Table

Delete	Frame Type	Value	Group Name
Delete	LLC	DSAP: 0xFF SSAP: 0xFF	

Group Name

A valid Group Name is a unique 16-character long string for every entry which consists of a combination of alphabets (a-z or A-Z) and integers (0-9).

Note: special character and underscore (_) are not allowed in the Group Name field.

Adding a New Group to VLAN mapping entry

Click the **Add New Entry** button to add a new entry in the mapping table. An empty row is added to the table; configure Frame Type, Value and the Group Name as needed. The **Reset** button can be used to undo the addition of a new entry.

Buttons

Save: Click to save changes.

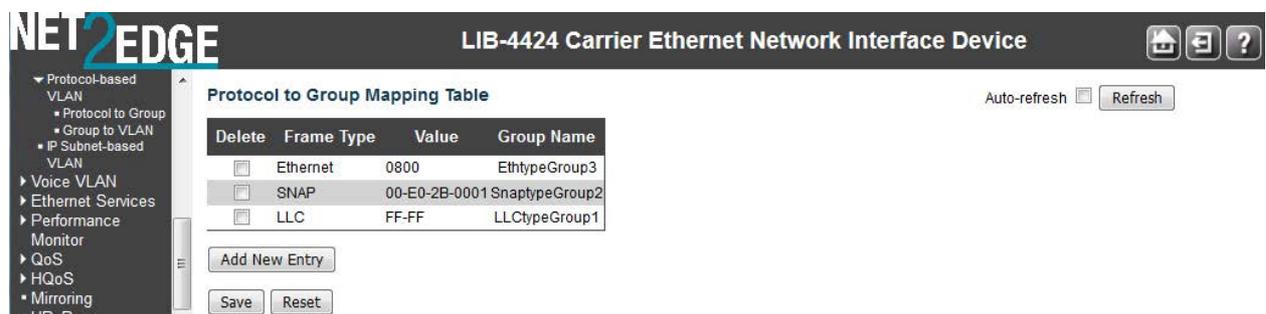
Reset: Click to undo any changes made locally and revert to previously saved values.

Auto-refresh: Check this box to refresh the page automatically every 3 seconds.

Refresh: Click to refresh the page immediately.

Example

The screen below shows three valid, saved Groups in the mapping table.

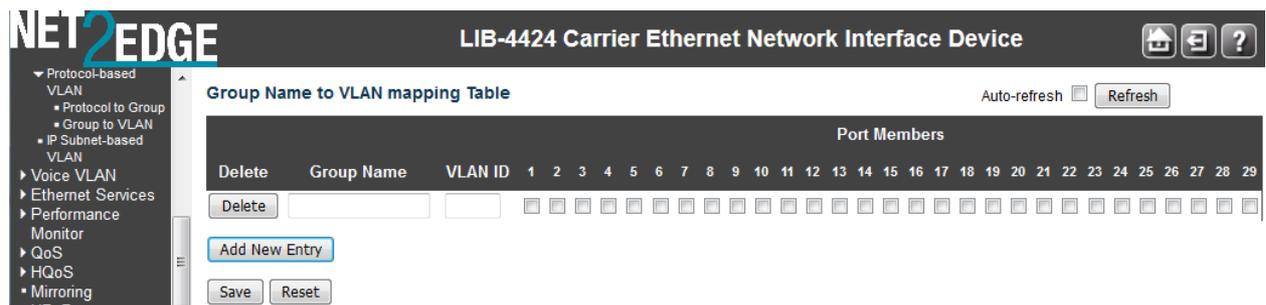


Group to VLAN

The **Configuration > VCL > Protocol-based VLAN > Group to VLAN** menu path lets you configure the LIB-44xx 'Group Name to VLAN mapping table' settings. This page lets you map an existing, configured Group Name to a VLAN for the selected switch.

At default, the table displays "No Group entries".

Click the **Add New Entry** button to display the entry fields.



The displayed settings are explained below:

Delete

To delete a Group Name to VLAN map entry, check this box. The entry will be deleted from the LIB-44xx at the next Save.

Group Name

A valid Group Name is a string of up to 16 characters which consists of a combination of up to 16 alpha (a-z or A-Z) and numeric (0-9) characters; no special characters are allowed. Whichever Group

name you try map to a VLAN must be present in the Protocol to Group mapping table (see above) and must not already be used by any other existing mapping entry on this page.

VLAN ID

Indicates the VID to which Group Name will be mapped. The valid VLAN ID range is 1-4094.

Port Members

A row of checkboxes for each port (e.g., ports 1-4 on the LIB-4400) is displayed for each Group Name to VLAN ID mapping. To include a port in a 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. At least one Port Member checkbox must be checked. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Add New Entry (Add a New Group to VLAN mapping table)

Click the **Add New Entry** button to add a new entry in the mapping table. An empty row is added to the table; configure the Group Name, VLAN ID and port members as needed. The valid VLAN ID values are **1** to **4094**.

The **Delete** button can be used to undo the addition of new entry.

Buttons

Auto-refresh: Check this box to refresh the page automatically. Automatic refresh occurs every 3 seconds.

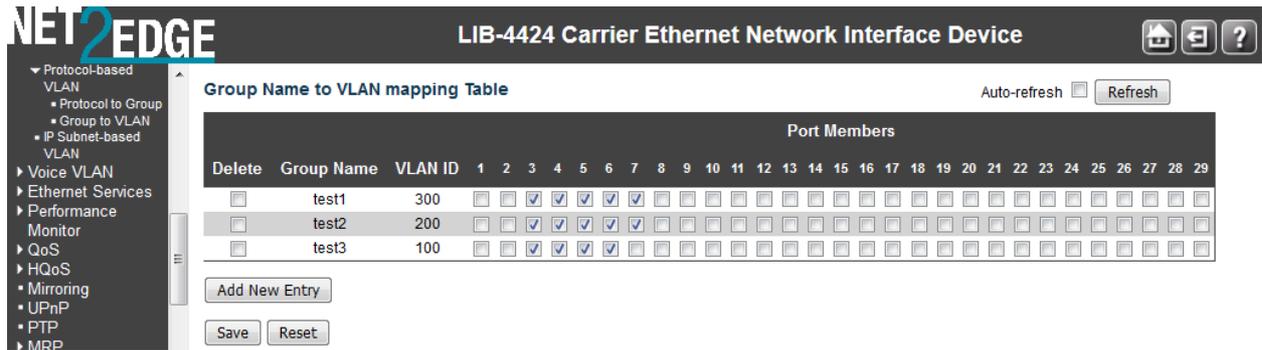
Refresh: Click to refresh the page immediately.

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Adding a New Group to VLAN mapping entry

The screen below shows three Group Names added to the Group Name to VLAN members table, as configured from the **Configuration - VCL - Group to VLAN** menu path.



Messages

Message: Invalid characters found. Please check help page for correct Group name format
Meaning: You entered one or more special characters in the **Group Name** field.

Recovery:

1. Click the **OK** button to clear the webpage message,
2. Re-enter the Group Name without any special characters (no space characters, underscore characters, dashes, etc.). You can enter up to 16 alpha (a-z or A-Z) characters and integers (0-9).
3. Click the **Save** button when done.

IP Subnet-based VLAN

IP subnet-based VLAN entries can be configured from the **Configuration > VCL > IP Subnet-based VLAN** menu path. With this method, a VLAN membership is defined by the subnet to which a device's IP address belongs.

This page lets you add, update, and delete IP subnet-based VLAN entries and assign the entries (membership) to different ports. VLANs are layer 2 constructs, compared with IP subnets which are layer 3 constructs. In an environment employing VLANs, a one-to-one relationship often exists between VLANs and IP subnets, but it is possible to have multiple subnets on one VLAN. VLANs and IP subnets provide independent Layer 2 and Layer 3 constructs that map to each other, and this correspondence helps in network design. Note that this page shows only static entries. Click the **Add New Entry** button to display the default page entry fields.

The screenshot shows the NET2EDGE web interface for a LIB-4424 Carrier Ethernet Network Interface Device. The main content area is titled "IP Subnet-based VLAN Membership Configuration". On the left is a navigation menu with options like "IP Subnet-based VLAN", "Voice VLAN", "Ethernet Services", "Performance", "Monitor", "QoS", "HQoS", "Mirroring", "UPnP", "PTP", "MRP", and "GVRP". The configuration area has a "Delete" checkbox, an "IP Address" field (0.0.0.0), a "Mask Length" field (24), and a "VLAN ID" field (1). Below these are "Add New Entry", "Save", and "Reset" buttons. To the right is a "Port Members" section with a row of 29 checkboxes, numbered 1 through 29.

The IP subnet-based VLAN entries are explained below:

Delete

To delete an IP subnet-based VLAN entry, check this box and click **Save**. The entry will be deleted from the switch.

VCE ID

Sets / shows the index of the entry. Its valid value is from **0- VCLIPIdMax**. If a VCE ID is **0**, the LIB-44xx will auto-generate the VCE ID for that entry. Deletion and lookup of IP subnet-based VLAN are based on VCE ID.

IP Address

Sets / shows the IP address. Enter a valid IP address in dotted decimal notation ('x.y.z.w') where x, y, and z are decimal numbers from 0 to 255.

Mask Length

Sets / shows the network mask length. The valid mask length range is 1-32.

VLAN ID

Sets / shows the VLAN ID. The VLAN ID can be changed for the existing entries.

Port Members

Displays a row of check boxes for each port for each IP subnet-based VLAN entry. To include a port in a IP subnet-based VLAN, check the checkbox (e.g., ports 1-4 on the LIB-4400). To remove or exclude the port from the IP subnet-based VLAN, make sure the box is unchecked. By default, no ports are members, and all boxes are unchecked. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Add New Entry (Add a New IP Subnet-based VLAN)

Click the **Add New Entry** button to add a new IP subnet-based VLAN entry. An empty row is added to the table, and the IP subnet-based VLAN entry can be configured as needed. Any IP address/mask can be configured for the IP subnet-based VLAN entry. Valid VLAN ID values are **1** through **4094**.

The IP subnet-based VLAN entry is enabled when you click the **Save** button. The **Delete** button can be used to undo the addition of new IP subnet-based VLANs.

Buttons

Auto-refresh: Check this checkbox to refresh the page automatically every three seconds.

Refresh: Refreshes the displayed table.

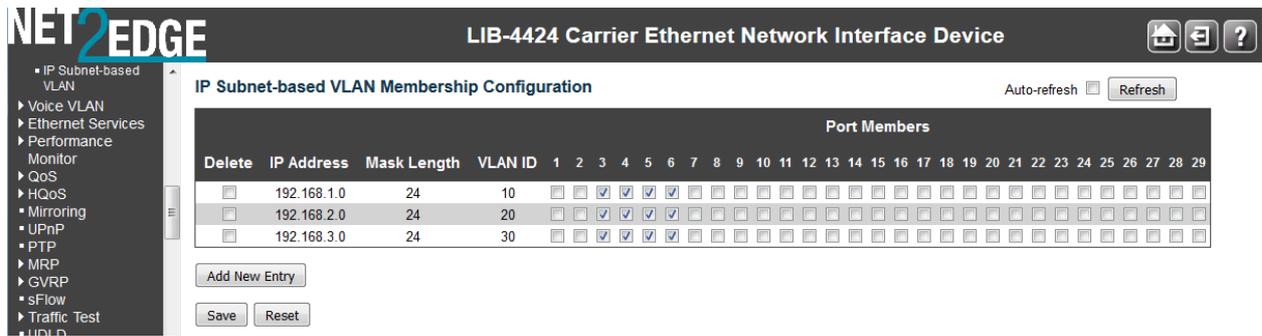
Add New Entry: Click to add a new IP subnet-based VLAN entry. See above.

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Example

In the screen sample below, three IP-subnet based VLAN VCE IDs have been created.



Each of the three IP-subnet based VLAN VCE IDs has a mix of IP addresses, Mask lengths, VLAN IDs, and port members assigned.

Ethernet Services Configuration

Ethernet services such as EPL, EVPL, E-LAN and E-Tree per MEF 6.2 will be supported in phases. The initial release supports EPL and EVPL services.

The EPL (Ethernet Private Line) is a point-to-point service which is totally transparent. To maintain transparency, the EPL service bundles all the customer traffic on an UNI into a tunnel (typically using provider S-tags) and tunnels all traffic in that SVLAN to the UNI on the other side.

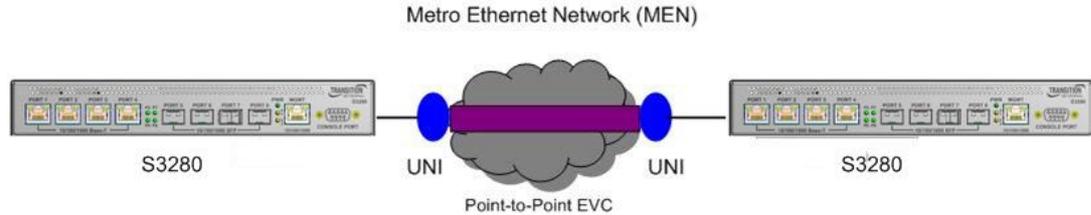


Figure 14a. EPL service per MEF 6.1

The EVPL service provides different services at a UNI, so transparency is not guaranteed like EPL service. Services from one UNI can reach different UNIs and not one-to-one as with EPL.

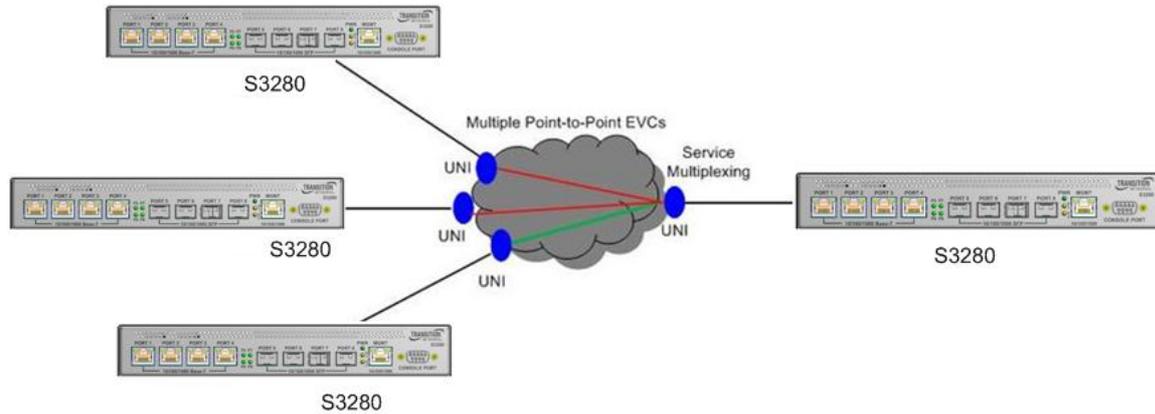


Figure 14b. EVPL Service per MEF 6.1

Bandwidth Profile (BWP)

A Bandwidth Profile (BWP) enforces bandwidth utilization according to the SLS (Service Level Specification) that the Subscriber and Service Provider (SP) have agreed on. The BWP can be viewed as the “enforcement of the long term average CIR (guaranteed bandwidth) and EIR (excess bandwidth) allowed by the service.

CIR (Committed Information Rate) is the bit rate for which the SP provides performance guarantees in terms of performance attributes for the service. EIR (Excess Information Rate) is the additional bit-rate that the subscriber can utilize and for which traffic can probably pass through the CEN as long as there is no congestion. Note that the total rate, sometimes called “PIR” (Peak Information Rate), is the sum of CIR and EIR.

The BWP classifies service frames into three ‘colours’, each of which denotes a specific compliance level:

Green – Frames within the CIR / CBS compliance level. These frames are subject to the SLS.

Yellow – Frames exceeding the CIR/CBS but are within the EIR/EBS. These frames are delivered as “best effort” but are not subject to the SLS. The CEN may drop some or all of these frames based on congestion conditions in the network and still be within the SLS.

Red – Frames not conforming to the BWP are dropped, either because the rate exceeds the sum of CIR + EIR, or because there are not enough yellow tokens to admit a frame that is within EIR/EBS. The two types of bandwidth profiles are Ingress BWP (which limits the rates of frames entering the CEN), and Egress BWP (which limits the rate of frames egressing the CEN, thus protecting overload towards the egress CE).

A BWP can be applied in three forms: 1) Per UNI, 2) Per EVC at a UNI, and 3) Per CoS ID per EVC at a UNI. These three types of ingress and egress bandwidth profiles can be applied at a UNI endpoint, and each provides increasingly more granularity in dividing the bandwidth at the UNI. Note that the algorithm used is the same for all three forms.

At a given UNI, at most one ingress BWP can be applied and at most one egress BWP can be applied to a given service frame (i.e., you cannot define an ingress BWP per UNI and an ingress BWP per EVC).

LIB-44xx BWP is configured from the **Configuration > Ethernet Services > Bandwidth Profiles** menu path.

Bandwidth Profile per UNI

The figure below shows an example of three EVCs sharing one BWP for the port. The bandwidth profile is controlled by configuring a dual leaky bucket policer for the entire port (UNI). Each EVC requires its own ASP to keep statistics specific to the EVC. Each EVC can have multiple Classes of Service; however, these are metered and counted separately per CoS.

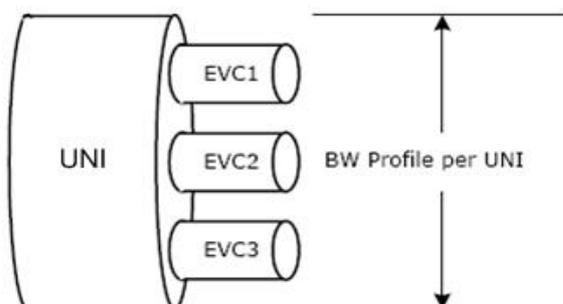


Figure 15. Bandwidth Profile per UNI

Bandwidth Profile per EVC

The figure below shows three EVCs, each having its own BWP. Each EVC requires its own ASP, and each ASP maps to its own Service Policer. Statistics are kept separately for each EVC. Each EVC can have multiple Classes of Service; however, these are metered and counted separately. EVCs that share one or more Classes of Service are metered and counted at the CoS level and not per EVC.

A port-level (UNI) DLB policer can be configured to control the bandwidth profile of the entire UNI on top of per EVC DLN policing (not shown). Bandwidth profiling at both the EVC level and the UNI level is enhanced as compared to MEF standards.

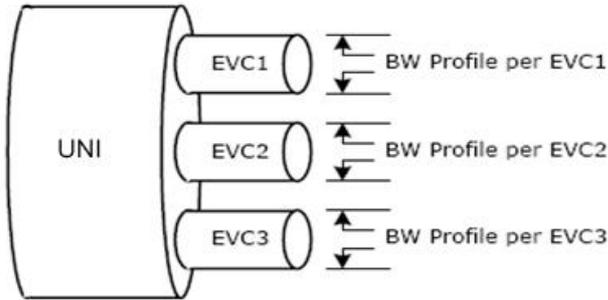


Figure 16. Bandwidth Profile per EVC

Bandwidth Profile per CoS ID per EVC

The figure below shows an example of an ingress BWP per CoS ID.

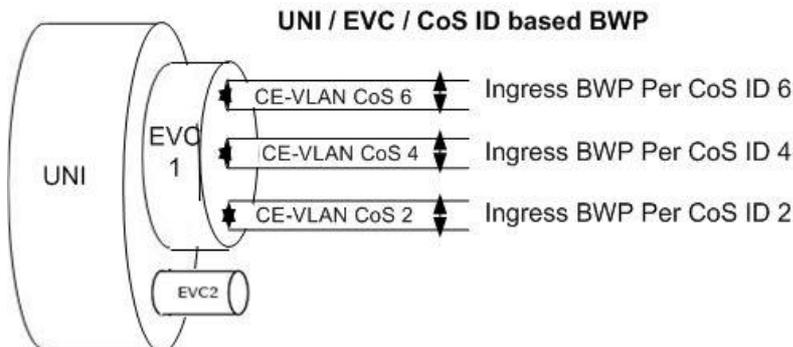


Figure 17. Bandwidth Profile per CoS ID per EVC

This bandwidth profile (Bandwidth Profile per CoS ID per EVC) enables the EVC bandwidth to be partitioned by CoS. The CoS ID is a set of one or more S-Tag PCP values. Each CoS ID indicates a single Class of Service instance. Frame colour can be marked in one of two methods: DEI bit or S-tag PCP values.

Configuration > Ethernet Services

From the **Configuration > Ethernet Services** menu path you can configure LIB-44xx Ethernet services in terms of ports, bandwidth profiles, EVCs (Ethernet Virtual Circuits) and ECEs (EVC Control Entries).

Configuration > Ethernet Services > Ports

This page lets you display and change current EVC port configuration settings.



The **Ethernet Services > Ports** configuration settings are explained below:

EVC ID

The EVC ID..

Role

The port role on the specific EVC. The possible values are:

Disabled: Not UNI/NNI.

NNI: NNI role.

Root: Root UNI role.

Leaf: Leaf UNI role.

Encapsulation

The encapsulation ID mapping on the specific EVC. The allowed range is from 0 through 907.

L2CP Profile

The L2CP Profile ID mapping on the specific EVC. The allowed range is from 0 through 62

HQoS

The HQoS ID mapping. The allowed range is from 0 through 255.

Cos ID Policer

The Cos ID Policer mapping on the specific EVC. For UNI ports, 8 policers are allocated (COSID 0-7), for NNI ports one policer is allocated (COSID 0).

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Configuration > Ethernet Services > Bandwidth Profiles

This page displays current EVC ingress bandwidth profile configurations. The policers configured here can be used to limit the traffic received on NNI ports. A policer can limit the bandwidth of received frames. Each policer is located in front of the ingress queue. The default BWP configuration page is shown below:

CoS ID	State	Type	Policer Mode	Rate Mode	CIR (kbps)	CBS (bytes)	EIR (kbps)	EBS (bytes)
0	Enabled	MEF	Blind	Data	10000	10000	50000	10000

The Ethernet Services ingress Bandwidth Profile Configuration settings are explained below:

COS ID

The CoS ID. For UNI ports, 8 policers are allocated (COSID 0-7), for NNI ports one policer is allocated (COSID 0)

State

The administrative state of the policer. The allowed values are:

Enabled: The policer enabled.

Disabled: The policer is disable.

Type

The type of the policer. The allowed values are:

MEF: MEF ingress policer.

Single: Single bucket policer policers.

Policer Mode

The colour mode of the bandwidth profile. The valid values are:

Coupled: Colour-aware mode with coupling enabled.

Aware: Colour-aware mode with coupling disabled.

Blind: Colour-blind mode (the colour disposition by a previous policer is NOT taken into account).

Rate Mode

The rate type of the policer. The allowed values are:

Data: Specify that this policer operates on data rate.

Line: Specify that this policer operates on line rate

CIR (kbps)

The Committed Information Rate of the bandwidth profile. The valid range is **0** to **1000000** kilobit per second. CIR (Committed Information Rate) is a Bandwidth Profile parameter that defines the average rate in bps of ingress Service Frames up to which the network delivers Service Frames and meets the performance objectives defined by the CoS Service Attribute. A Bandwidth Profile property, where a pre-determined level of Bandwidth Profile compliance for each Service Frame, if present, is ignored when determining the level of compliance for each Service Frame.

CBS (bytes)

The Committed Burst Size of the bandwidth profile. The valid range is **0** to **100000** bytes. CBS (Committed Burst Size) is a Bandwidth Profile parameter that limits the maximum number of bytes available for a burst of Service Frames sent at the UNI speed to remain CIR-conformant. It defines the average rate in bps of ingress Service Frames up to which the network delivers Service Frames and meets the performance objectives (as defined by the CoS Service Attribute).

EIR (kbps)

The Excess Information Rate of the bandwidth profile. The valid range is **0** to **1000000** kilobit per second.

EIR (Excess Information Rate) is a Bandwidth Profile parameter that defines the average rate in bps of ingress Service Frames up to which the network may deliver Service Frames without any performance objectives. This is the rate up to which the network will attempt to deliver Ethernet service frames before they are discarded.

EBS (bytes)

The Excess Burst Size of the bandwidth profile. The valid range is **0** to **100000** bytes. EBS (Excess Burst Size) is the maximum number of bytes allowed for incoming Service Frames to be EIR-conformant (yellow frames).

Buttons

Refresh: Refreshes the displayed table starting from the input fields.

|<<: Updates the table, starting with the first entry in the table.

<<: Updates the table, ending at the entry before the first entry currently displayed.

>>: Updates the table, starting with the entry after the last entry currently displayed.

>>|: Updates the table, ending at the last entry in the table.

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Ethernet Services EVCs

The **Configuration > Ethernet Services > EVCs** menu path displays current EVC configurations.

The EVC settings can also be configured here. On this system, only Provider Bridge based EVCs are supported.

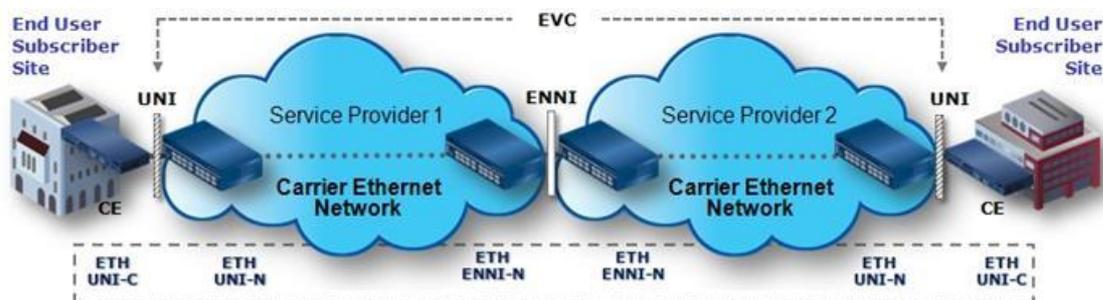


Figure 18. Provider Bridge E-LINE Service

The EVC (Ethernet Virtual Connection) is an association of two or more UNIs that limits the exchange of frames to UNIs in the Ethernet Virtual Connection. The User Network Interface (UNI) is the physical interface or port that is the demarcation between the customer and the service provider/Cable Operator/Carrier/MSO. The UNI is the physical demarcation point between the responsibility of the Service Provider and the responsibility of the Subscriber.

MEF Ethernet Virtual Connection Types

The Metro Ethernet Forum (MEF) specifies these EVC (Ethernet Virtual Connection) types:

E-Line EVC: Point-to-point Service Ethernet Private Line (EPL) allows only one EVC per UNI port, while Ethernet Virtual Private Line (EVPL) allows multiple EVCs per UNI port. The figure below shows a Provider Bridge E-LINE service example.

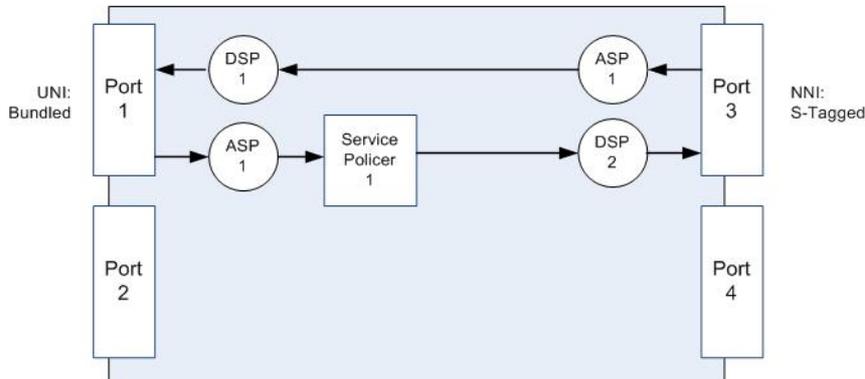


Figure 19. Provider Bridge E-LINE Service

E-LAN EVC: Multipoint Service Ethernet Private LAN (EP-LAN) allows only one EVC per UNI port, while Ethernet Virtual Private LAN (EVP-LAN) allows multiple EVCs per UNI port. This is a bridged service.

E-TREE EVC: Rooted Multipoint Service Ethernet Private Tree (EP-TREE) allows only one EVC per UNI port, while Ethernet Virtual Private Tree (EVP-TREE) allows multiple EVCs per UNI port. This is a bridged service where the root port has access to all the leaf ports, but the leaf ports only have access to the root port.

Configuration Prerequisites

E-LINE is supported; you must disable MAC Learning first for E-LINE support.

E-LAN is supported; you must remove each VLAN and enable MAC Learning for E-LAN support.

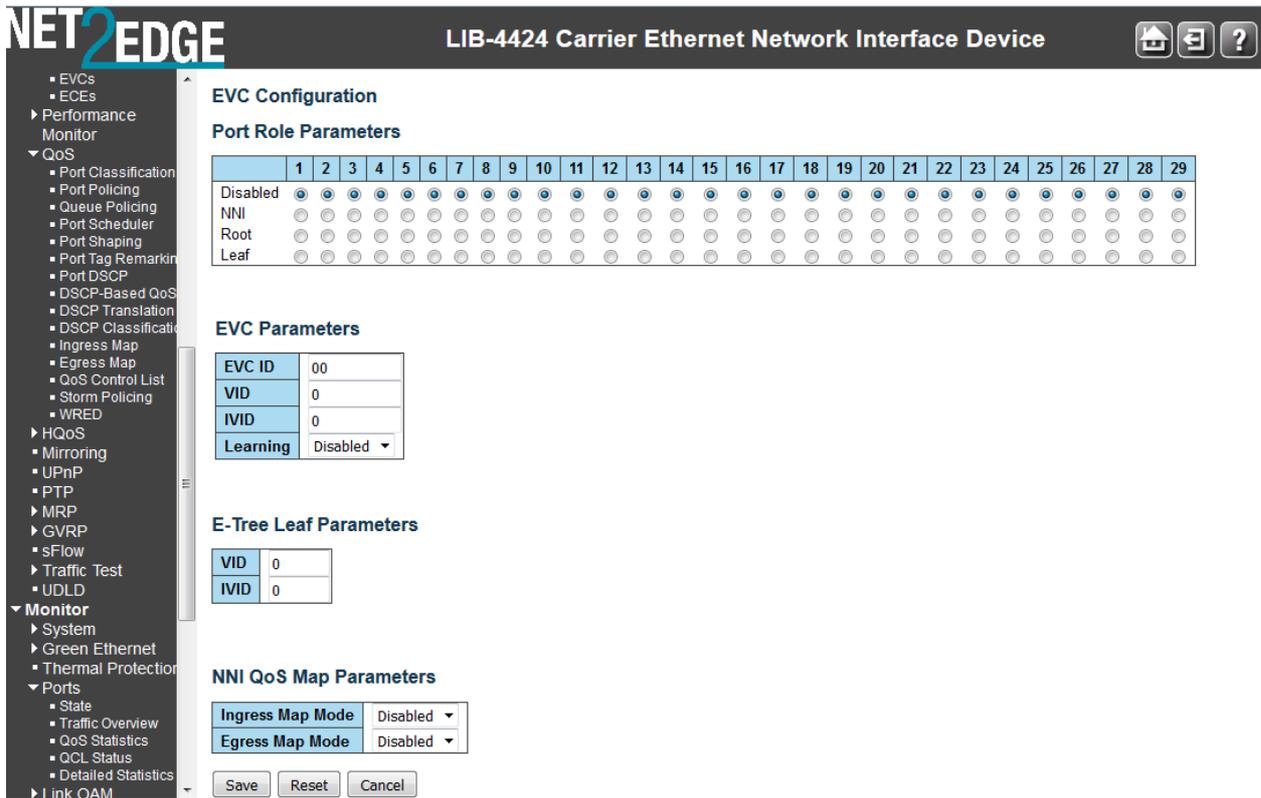
Create the EVC first, and then create the ECE next.

Configuration > Ethernet Services > EVCs

This menu path displays the EVC Control List Configuration table. You can add, edit, and delete EVCs here. The EVC configuration parameters on the default page are shown below:

EVC ID	VID	IVID	Learning	Port Role	Leaf				NNI QoS Map		HQoS IDs
					VID	IVID	Ingress	Egress	Ingress	Egress	
1	1000	1000	Disabled	NNI:5,6	0	0	Disabled	Disabled	None	Configure	

Click the plus sign (+) modification button to add a new EVC. The EVC Configuration page displays.



Configure the new EVC’s NNI Ports, EVC Parameters, Inner Tag and Outer Tag as explained below:

NNI Ports

The list of Network to Network Interfaces for the EVC (checkboxes for ports 1-4 on the LIB-4400). Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

EVC ID

The EVC ID identifies the EVC. The valid range is **1 - 128**. This is the EVC ID for the ECE. The ECE is only active when mapping to an existing EVC.

VID

The VLAN ID in the Provider Bridge network. It may be inserted in a C-tag, S-tag, or S-custom tag depending on the NNI port VLAN configuration. The valid range is **1 - 4094**.

IVID

The Internal/classified VLAN ID in the Provider Bridge network. The valid range is **1 - 4094**. The IVID is different from the default VLAN and other VLANs used for normal forwarding. This IVID is added to the UNI/NNI ports involved in the EVC.

The IVID parameter determines the internal VLAN ID assigned to the frames mapping to the EVC rule. Frame forwarding is limited to the port members configured for the IVID. If learning is enabled, source addresses are learned on the IVID.

The EVC control module automatically adds UNI/NNI ports as VLAN members for the IVID. The purpose of the parameter is to allow using one VID inserted in tags on NNI ports and using another VID locally in the LIB-44xx. The IVID can be used to separate the local VLAN space from the VLAN

space of the rest of the network. In some cases, only a few internal VLANs are needed to support a bigger number of external VLANs.

Learning

The learning mode for the EVC controls whether source MAC addresses are learned for frames matching the EVC. Learning may be disabled if the EVC only includes two UNI/NNI ports. The valid values are:

Enabled: Learning is enabled (MAC addresses are learned).

Disabled: Learning is disabled (MAC addresses are not learned).

Policer ID Filter

The ingress bandwidth profile mode for the EVC. The possible values are:

Specific: Lets you configure a specific policer. The valid range is from **1** - **128**.

Discard: All received frames are discarded for the EVC.

None: No bandwidth profile for the EVC.

Policer ID Value

The policer ID specific value for the EVC (e.g., policer ID **1**). Only applies if the “Policer ID Filter” setting above is set to “Specific”. The valid range is from **1** - **128**.

Click the **Save** button when done configuring.

Modification Buttons

You can modify each EVC in the table using the following buttons:

: Edits the current EVC row. Displays the EVC Configuration page for the row.

: Deletes the EVC.

: Adds a new EVC. Displays the EVC Configuration page for the new row.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Cancel: Return to the previous page; any changes made locally will be undone.

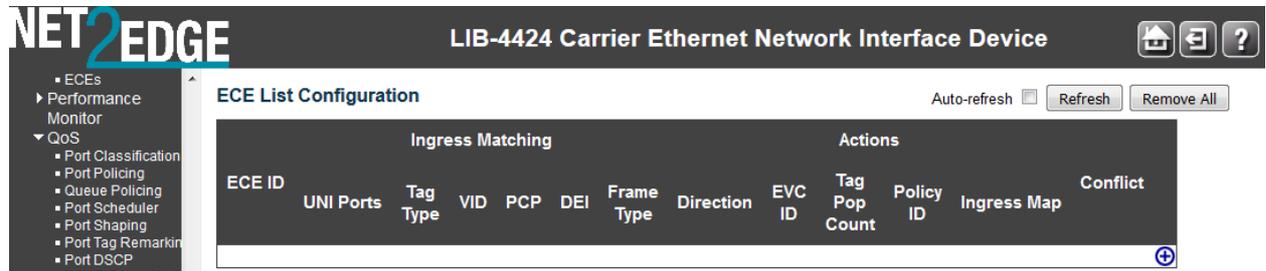
Example

A sample EVC Control List Configuration is shown below with two EVCs (EVC ID 1 and 2).

Click the modification button to add one or more additional new EVCs.

Configuration > Ethernet Services > ECEs

This menu path displays the EVC Control Entries (ECEs). You can add, edit, and delete ECEs here. Make sure that you create the EVCs first (previous section), and then create the ECEs next. The default ECE Control List Configuration page is shown below:



From the default page, you click the plus sign (+) modification button in the far right column to add a new entry to the ECE listing at the ECE Configuration page.

The ECE Control List Configuration parameters are explained below:

ECE ID

The ECE ID identifies the ECE. Unique ECE IDs are automatically assigned to ECEs added. The valid range is **1** to **128**.

UNI Matching

UNI Ports

The list of User Network Interfaces for the ECE (e.g., IETF, MEF). The physical interface or port that is the demarcation between the customer and the service provider/Cable Operator/Carrier/MSO.

Tag Type

The tag type for the ECE. The possible values are:

Tagged: The ECE will match tagged frames only.

Untagged: The ECE will match untagged frames only.

Any: The ECE will match both tagged and untagged frames.

undefined: The ECE tag type is not known or is unfamiliar.

VID

The VLAN ID for the ECE. The VLAN ID value is only significant if tag type 'Tagged' is selected. Valid values are:

Specific: The valid range is **1** to **4094**.

Any: The ECE will match any VLAN ID.

PCP

The PCP value for the ECE. The PCP value is only significant if tag type 'Tagged' is selected. The Priority Code Point (PCP) is a 3-bit field storing the priority level for the 802.1Q frame (also known as User Priority). Valid values are:

Specific: The ECE will match a specific PCP in the range **0** through **7**.

Range: The ECE will match PCP values in the selected range **0-1**, **2-3**, **4-5**, **6-7**, **0-3** or **4-7**.

Any: The ECE will match any PCP value.

DEI

The DEI value for the ECE. The DEI value is only significant if tag type 'Tagged' is selected. The Drop Eligible Indicator (DEI) is a 1-bit field in the VLAN tag. The valid values are **0**, **1** or **Any**.

Frame Type

The frame type for the ECE. This selection defines some of the additional fields to be displayed. The valid values are:

Any: The ECE will match any frame type.

IPv4: The ECE will match IPv4 frames only.

IPv6: The ECE will match IPv6 frames only.

Actions

Direction

The EVCs and ECEs are used to setup flows in one or both directions as determined by the ECE Direction parameter. If the ECE is bidirectional, the ingress rules of the NNI ports will be setup to match the traffic being forwarded to NNI ports. The possible values are:

Both: Bidirectional.

UNI-to-NNI: Unidirectional from UNI to NNI.

NNI-to-UNI: Unidirectional from NNI to UNI. The feature "NNI Outer Tag" on the ECE Configuration page is ONLY applied when the Actions - Direction is set to "NNI-to-UNI". The 'NNI-to-UNI Tag Mode' is only used when the direction is 'NNI-to-UNI'. It only applies to the 'NNI-to-UNI' direction. The 'NNI-to-UNI Tag Mode' selection is disabled when the direction is not 'NNI-to-UNI'.

The 'NNI-to-UNI Tag Mode' may be used in situations where you would like to translate from C-tag to S-tag and optionally also change the VLAN ID. Frames traversing from NNI-to-UNI will pop an S-tag, and on egress from the UNI push a new C-tag with the VLAN ID from '**EVC Configuration > Outer Tag > VLAN ID**'. In the other direction you should be able to do similar actions in the ECE using '**Tag Pop Count**'.

Note: Two unidirectional services in each direction do not equate to a bidirectional service. Also, it is recommended to make the NNI ports as C/S-ports and to add VLAN membership entries accordingly to ensure proper operation on the EVCs.

EVC ID

The EVC ID for the ECE. The ECE is only active when mapping to an existing EVC. Valid values are:

Specific: The valid range is **1** - **128**.

None: The ECE does not map to an EVC (to allow for per-CoS bandwidth profile).

Policer ID

The ingress bandwidth profile mode for the ECE. The possible values are:

Specific: The range is from **1** - **128**.

Discard: All received frames are discarded for the ECE.

None: All received frames are forwarded for the ECE.

Tag Pop Count

The ingress tag pop count for the ECE. The valid range is from **0** - **2**. Each service instance can change the existing VLAN tag to a new VLAN tag by adding, removing, or translating one or two VLAN tags.

A 'flexible' VLAN tag rewrite can include three main operations: Pop - pop the tag (remove an existing tag), or Push (add a new tag) or Translate (change one or two tags to another one or two tags - a combination of pop and push operations).

Policy ID

The ACL Policy ID for the ECE for matching ACL rules. Valid values are **0** - **255**.

NNI Outer Tag

NNI-to-UNI Tag Mode

The outer tag for NNI-to-UNI direction for the ECE. The feature "NNI Outer Tag" on the ECE Configuration page is ONLY applied when the Actions - Direction is set to "NNI-to-UNI".

The valid values are:

Enabled: Enable outer tag for NN-to-UNI direction for the ECE.

Disabled: Disable outer tag for NNI-to-UNI direction for the ECE.

The 'NNI-to-UNI Tag Mode' can only be enabled when the direction is 'NNI-to-UNI'. It only applies to the 'NNI-to-UNI' direction. The 'NNI-to-UNI Tag Mode' selection is disabled when the direction is not 'NNI-to-UNI'. The 'NNI-to-UNI Tag Mode' may be used in situations where you would like to translate from C-tag to S-tag and optionally also change the VLAN ID. Frames traversing from NNI-to-UNI will pop an S-tag, and on egress from the UNI push a new C-tag with the VLAN ID from '**EVC Configuration > Outer Tag > VLAN ID**'. In the other direction you should be able to do similar actions in the ECE using '**Tag Pop Count**'.

Outer Tag VID

The EVC outer tag VID for UNI ports. The valid range is **0** - **4094**.

Tag PCP/DEI Preservation

The outer tag PCP and DEI preservation for the ECE. The possible values are:

Preserved: The outer tag PCP and DEI are preserved.

Disable: The outer tag PCP and DEI are fixed.

Outer Tag PCP

The outer tag PCP (Priority Code Point) value for the ECE. The valid range is **0** to **7**.

Outer Tag DEI

The outer tag DEI value for the ECE. The valid values are **0** or **1 or undefined** (the ECE tag type is not known or is unfamiliar).

Conflict

Indicates the hardware status of the specific ECE. The specific ECE is not applied to the hardware due to hardware limitations.

NNI Inner Tag

Inner Tag Type

The inner type for the ECE determines whether an inner tag is inserted in frames forwarded to NNI ports. The possible values are:

None: An inner tag is not inserted.

C-tag: An inner C-tag is inserted.

S-tag: An inner S-tag is inserted.

S-custom-tag: An inner tag is inserted and the tag type is determined by the VLAN port configuration of the NNI.

Inner Tag VLAN ID

The inner tag VLAN ID for the ECE. The allowed range is from **0** through **4094**.

Inner Tag PCP/DEI Preservation

The inner tag PCP and DEI preservation for the ECE. The possible values are:

Preserved: The inner tag PCP and DEI is preserved.

Fixed: The inner tag PCP and DEI is fixed.

Inner Tag PCP

The inner tag PCP value for the ECE. The allowed range is from **0** through **7**.

Inner Tag DEI

The inner tag DEI value for the ECE. The allowed value is **0** or **1**.

Modification Buttons

You can modify each ECE (EVC Control Entry) in the table using the following buttons:

- : Inserts a new ECE before the current row.
- : Edits the ECE row.
- : Moves the ECE up the list.
- : Moves the ECE down the list.
- : Deletes the ECE.
- : The lowest plus sign adds a new entry at the bottom of the ECE listings.

Buttons

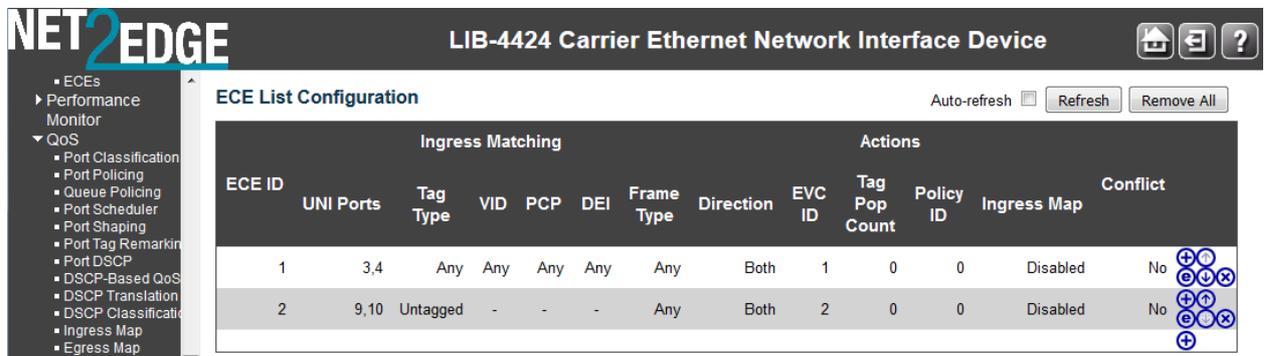
Auto-refresh: Check the checkbox to refresh the page automatically every three seconds.

Refresh: Click to refresh the page.

Remove All: Click to remove all ECEs.

Example

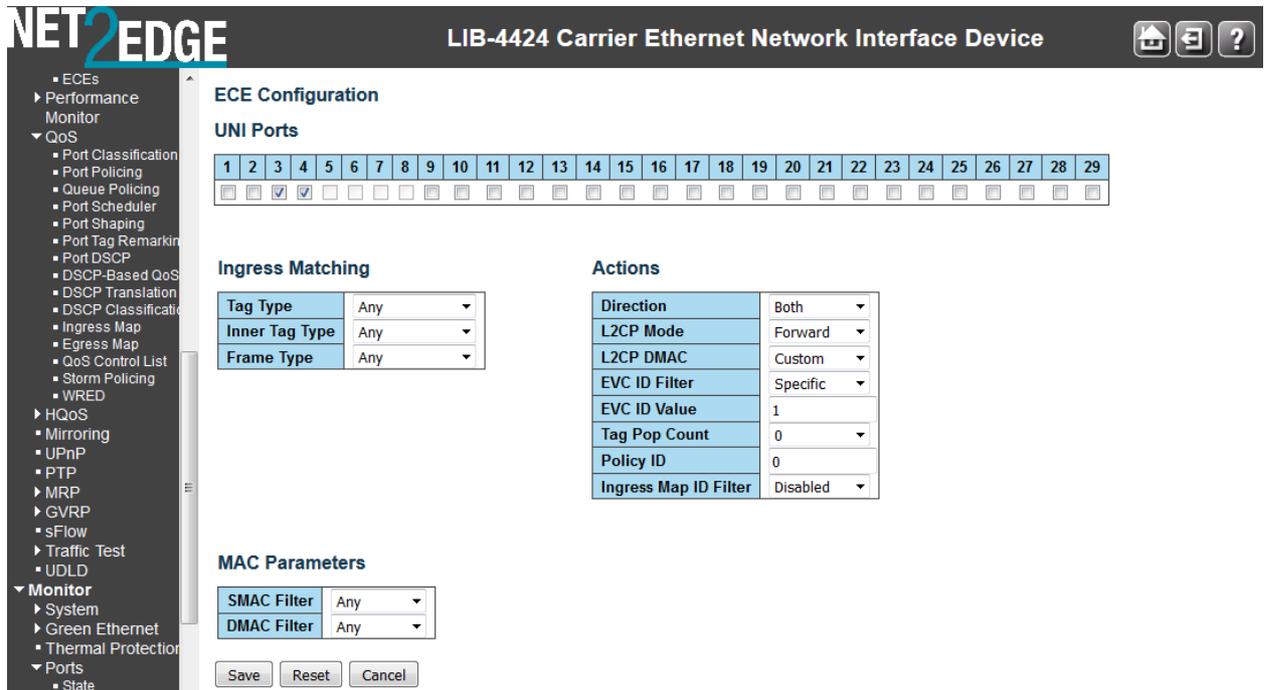
The screen below shows two valid, saved ECEs.



Use the Modification buttons to add, edit, delete, or move ECE instances in the table.

ECE Configuration Page

With existing entries, when you click one of the plus sign (+) modification buttons to add a new entry to the ECE listings, the **ECE Configuration** page displays.



Configure the new ECE’s UNI Ports, UNI Matching, Actions, NNI Outer Tag, and/or NNI Inner Tag parameters as explained above. **Note:** the set of parameters displayed here depends on the **Frame Type** selection at the **UNI Matching** section.

The **ECE Configuration** page parameters are explained below:

UNI Ports

The list of User Network Interfaces for the ECE. Check or uncheck one or more of the checkboxes (e.g., ports 1-4 on the LIB-4400). Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

UNI Matching

Tag Type

The tag type for matching the ECE. The valid values are:

Any: The ECE will match both tagged and untagged frames.

Untagged: The ECE will match untagged frames only.

Tagged: The ECE will match tagged frames only.

VLAN ID Filter

The **VLAN** ID filter for matching the ECE. It only significant if tag type 'Tagged' is selected. The possible values are:

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 value with this ECE, choose this value. A field for entering a specific value displays.

Range: If you want to filter a specific VLAN ID range filter with this ECE, choose this value. A field for entering a range displays.

VLAN ID Value

When "Specific" is selected for the VLAN ID filter, you can enter a specific value. The valid range is **0 - 4094**.

VLAN ID Range

When "Range" is selected for the VLAN ID filter, you can enter a specific range. The valid range is **0 - 4095**.

PCP

The PCP value for matching the ECE. It only significant if tag type 'Tagged' is selected. The valid values are:

Any: The ECE will match any PCP value.

Specific: The ECE will match a specific PCP in the range **0** through **7**.

Range: The ECE will match PCP values in the selected range **0-1**, **2-3**, **4-5**, **6-7**, **0-3** or **4-7**.

DEI

The DEI value for matching the ECE. It only significant if tag type 'Tagged' is selected. The valid values are **0**, **1** or **Any**.

Inner Tag Type

The inner tag type for matching the ECE. The valid values are:

Any: The ECE will match both tagged and untagged frames.

Untagged: The ECE will match untagged frames only.

Tagged: The ECE will match tagged frames only.

Inner VLAN ID Filter

The inner VLAN ID filter for matching the ECE. It only significant if Inner Tag Type 'Tagged' is selected. The valid values are:

Any: No inner VLAN ID filter is specified. (Inner VLAN ID filter status is "don't-care".)

Specific: If you want to filter a specific inner VLAN ID value with this ECE, choose this value. A field for entering a specific value displays.

Range: If you want to filter a specific inner VLAN ID range filter with this ECE, choose this value. A field for entering a range displays.

UNI Matching

Tag Type	Tagged
VLAN ID Filter	Range
VLAN ID Range	1 - 4095
PCP	Any
DEI	Any
Inner Tag Type	Tagged
Inner VLAN ID Filter	Specific
Inner VLAN ID Value	1
Inner PCP	1
Inner DEI	0
Frame Type	IPv6
DSCP Filter	Range
DSCP Range	0 - 63

Inner Tag VLAN ID Range

The range setting need based on bit mask concept. For example, **4-5** (last bit mask), **4-7** (last two bits mask). This parameter displays only if "Range" was selected in the "Inner VLAN ID Filter" field above.

Inner VLAN ID Value

When "Specific" is selected for the VLAN ID filter, you can enter a specific value. Valid values are **0 - 4094**.

Inner VLAN ID Range

When "Range" is selected for the VLAN ID filter, you can enter a specific range. Valid values are **0 - 4094**.

Inner Tag PCP

The inner PCP value for matching the ECE. It only significant if inner tag type 'Tagged' is selected. The possible values are:

Any: The ECE will match any PCP value.

Range: The ECE will match PCP values in the selected range **0-1, 2-3, 4-5, 6-7, 0-3** or **4-7**.

Specific: The ECE will match a specific PCP in the range **0** through **7**.

Inner Tag DEI

The inner DEI value for matching the ECE. It only significant if inner tag type 'Tagged' is selected. The valid values are **0, 1** or **Any**.

Frame Type

The frame type for the ECE. The possible values are:

Any: The ECE will match any frame type.

IPv4: The ECE will match IPv4 frames only.

IPv6: The ECE will match IPv6 frames only.

DSCP Filter

The DSCP filter for matching the ECE. The possible values are:

Any: No DSCP filter is specified. (DSCP filter status is "don't-care".)

Specific: If you want to filter a specific DSCP value with this ECE, choose this value. A field for entering a specific value appears.

Range: If you want to filter a specific DSCP range filter with this ECE, choose this value. A field for entering a range appears.

DSCP Value

When "Specific" is selected for the DSCP filter, you can enter a specific value. Valid values are **0 - 63**.

DSCP Range

When "Range" is selected for the DSCP filter, you can enter a specific range. Valid values are **0 - 63**.

Actions

Direction

The EVCs and ECEs are used to setup flows in one or both directions as determined by the ECE Direction parameter. If the ECE is bidirectional, the ingress rules of the NNI ports will be setup to match the traffic being forwarded to NNI ports. The valid values are:

Both: Bidirectional.

UNI-to-NNI: Unidirectional from UNI to NNI.

NNI-to-UNI: Unidirectional from NNI to UNI. The 'NNI-to-UNI Tag Mode' is only used when the direction is 'NNI-to-UNI'. It only applies to the 'NNI-to-UNI' direction. The 'NNI-to-UNI Tag Mode' selection is disabled when the direction is not 'NNI-to-UNI'.

The 'NNI-to-UNI Tag Mode' may be used in situations where you would like to translate from C-tag to S-tag and optionally also change

Actions

Direction	NNI-to-UNI <input type="button" value="v"/>
EVC ID Filter	Specific <input type="button" value="v"/>
EVC ID Value	1
Policer ID Filter	Specific <input type="button" value="v"/>
Policer ID Value	0
Tag Pop Count	1 <input type="button" value="v"/>
Policy ID	0

the VLAN ID. Frames traversing from NNI-to-UNI will pop an S-tag, and on egress from the UNI push a new C-tag with the VLAN ID from '**EVC Configuration > Outer Tag > VLAN ID**'. In the other direction you should be able to do similar actions in the ECE using '**Tag Pop Count**'.

Note: Two unidirectional services in each direction do not equate to a bidirectional service. Also, it is recommended to make the NNI ports as C/S-ports and to add VLAN membership entries accordingly to ensure proper operation on the EVCs.

EVC ID Filter

The EVC ID for the ECE. The ECE is only active when mapping to an existing EVC. The valid values are:

Any: No EVC ID filter is specified. (EVC ID filter status is "don't-care".)

Specific: If you want to filter a specific EVC ID with this ECE, choose this value. A field for entering a specific value displays.

EVC ID Value

When "Specific" is selected for the VLAN ID filter, you can enter a specific value. The valid range is **1 - 128**.

Policer ID Filter

The policer ID filter for matching the ECE. The valid values are:

Specific: If you want to filter a specific policer ID value with this ECE, choose this value. The "Policer ID Value" field displays for entering a specific value. Enter a specific value (**1 - 128**).

Discard: All received frames are discarded for the ECE.

None: All received frames are forwarded for the ECE. The bandwidth profile for the specified EVC ID is used.

EVC: The bandwidth profile for the specified EVC ID is used.

Tag Pop Count

The ingress tag pop count for the ECE (where "pop" indicates that the outer VLAN tag of the incoming frame is removed). The valid range is **0 - 2**.

Policy ID

The ACL Policy ID for the ECE for matching ACL rules. The valid range is **0 - 255**.

Class

The traffic class for the ECE. Valid values are **0 - 8** or **disabled**.

NNI Outer Tag

NNI-to-UNI Tag Mode

The outer tag for NNI-to-UNI direction for the ECE. The feature "NNI Outer Tag" on the ECE Configuration page is **only** applied when the Actions - Direction is set to "NNI-to-UNI". The valid values are:

Enable: Enable outer tag for NNI-to-UNI direction for the ECE.

Disable: Disable outer tag for NNI-to-UNI direction for the ECE.

The 'NNI-to-UNI Tag Mode' is only used when the direction is 'NNI-to-UNI'. It only applies to the 'NNI-to-UNI' direction. The 'NNI-to-UNI Tag Mode' selection is disabled when the direction is not 'NNI-to-UNI'.

The 'NNI-to-UNI Tag Mode' may be used in situations where you would like to translate from C-tag to S-tag and optionally also change the VLAN ID. Frames traversing from NNI-to-UNI will pop an S-

NNI Outer Tag

NNI-to-UNI Tag Mode	Disabled ▾
VLAN ID	1
PCP/DEI Preservation	Fixed ▾
PCP	0 ▾
DEI	0 ▾

tag, and on egress from the UNI push a new C-tag with the VLAN ID from '**EVC Configuration** > **Outer Tag** > **VLAN ID**'. In the other direction you should be able to do similar actions in the ECE using '**Tag Pop Count**'.

Outer Tag PCP/DEI Preservation

The outer tag PCP and DEI preservation for the ECE. The valid values are:

Preserved: The outer tag PCP and DEI is preserved.

Fixed: The outer tag PCP and DEI is fixed.

Outer Tag PCP

The outer tag PCP value for the ECE. The valid range is **0** to **7**.

Outer Tag DEI

The outer tag DEI value for the ECE. Valid values are **0** or **1**.

NNI Inner Tag

Type (NNI Inner Tag Type)

None: An inner tag is not inserted.

C-tag: An inner C-tag is inserted.

S-tag: An inner S-tag is inserted.

S-custom-tag: An inner tag is inserted and the tag type is determined by the VLAN port configuration of the NNI.

NNI Inner Tag

Type	None
VLAN ID	1
PCP/DEI Preservation	Fixed
PCP	0
DEI	0

VLAN ID (NNI Inner Tag)

The inner tag VLA ID (VID) for the VCE. The valid range is 1-4094.

PCP/DEI Preservation (NNI Inner Tag)

The inner tag PCP and DEI preservation for the ECE. The valid values are:

Preserved: The inner tag PCP and DEI is preserved.

Fixed: The inner tag PCP and DEI is fixed.

PCP (NNI Inner Tag)

The inner tag PCP (Priority Code Point) value for the ECE. Valid values are **0-7**.

DEI (NNI Inner Tag)

The inner tag DEI (Drop Eligible Indicator) value for the ECE. Valid values are **0** or **1**.

IPv6 Parameters

These parameters display only if IPv6 is selected as the **Frame Type** selection in the **UNI Matching** section.

Protocol

Any: Use any of the valid IPv6 protocols (UDP, TCP, or Other). This is the default setting.

UDP: Use only UDP as the valid IPv6 protocol.

TCP: Use only TCP as the valid IPv6 protocol.

Other: Select another Ipv6 protocol value other than Any, UDP, or TCP. When selected, the "Protocol Value" field displays (see below).

SIP/DIP Filter

Where Filter = process all flows before the next step, based on:

Any: Use any of the valid SIP/DIP (source IP / destination IP) filters. Where Filter = process all flows before the next step).

Specific: Specify a specific valid SIP/DIP (source IP / destination IP) filter to be used.

When "IPv6" is selected for the Frame Type, the field only supports 32 bits for IPv6 address mask.

DSCP Filter

Any: Use any of the valid DSCP filters.

Specific: Specify a specific valid DSCP filter to be used.

Range: Specify a series of contiguous valid DSCP filters to be used (0-63). If selected, you must also specify a "SIP/DIP Address" and a "SIP/DIP Mask" (below).

DSCP Value

When "Specific" is selected for the DSCP filter, you can enter a specific value. The valid value is 0 - 63.

DSCP Range

When "Range" is selected for the DSCP filter, you can enter a specific range. The valid range is 0 - 63.

SIP/DIP Address

Enter up to 32 bits following the "0x" prefix. This selection only displays if you selected "Range" as the "DSCP Filter" selection (above). This is the address to be used for the SIP/DIP (source IP / destination IP).

Enter an 8-character address from 00000000 - ffffffff.

SIP/DIP Mask

Enter up to 32 bits following the "0x" prefix. This selection only displays if you selected "Range" as the "DSCP Filter" selection (above). This is the mask to be used for the SIP/DIP (source IP / destination IP).

Enter an 8-character address from 0 - ffffffff.

Source Port Filter

Any: Use any of the valid source port filters.

Specific: Specify a specific valid source port filter to be used.

Range: Specify a series of contiguous valid source port filters to be used.

Dest. Port Filter

Any: Use any of the valid destination port filters.

Specific: Specify a specific valid destination port filter to be used.

Range: Specify a series of contiguous valid destination port filters to be used.

Protocol Value

Select the IPv6 protocol value to be used. The default is 0. The valid range is 0-255. This field displays only if "Other" is selected in the IPv6 "Protocol" field above.

IPv4 Parameters

These parameters display only if IPv4 is selected as the **Frame Type** selection in the **UNI Matching** section.

Protocol

The IP protocol for matching the ECE. The valid values are:

Any: No protocol filter is specified. (Protocol filter status is "don't-care".)

UDP: Specify the UDP for matching the ECE.

TCP: Specify the TCP for matching the ECE.

Other: Specify another protocol value for matching the ECE. When selected, the "Protocol Value" field displays (see below).

SIP/DIP Filter

Where Filter = process all flows before the next step, based on:

Any: Use any of the valid SIP/DIP (source IP / destination IP) filters.

Host: Use the host device's SIP/DIP (source IP / destination IP) filter as the basis.

Network: Use the network's SIP/DIP (source IP / destination IP) filter as the basis.

SIP/DIP Address

When "Host" or "Network" is selected for the SIP/DIP Filter, you can enter a specific host or network address.

SIP/DIP Mask

When "Host" or "Network" is selected for the SIP/DIP Filter, you can enter a specific network mask.

DSCP Filter

The DSCP filter for matching the ECE. The valid values are:

Any: No DSCP filter is specified. (DSCP filter status is "don't-care".)

Specific: If you want to filter a specific DSCP value with this ECE, choose this value. A field for entering a specific value displays.

Range: If you want to filter a specific DSCP range filter with this ECE, choose this value. A field for entering a range displays.

DSCP Value

When "Specific" is selected for the DSCP filter, you can enter a specific value. The valid value is **0 - 63**.

DSCP Range

When "Range" is selected for the DSCP filter, you can enter a specific range. The valid value is **0 - 63**.

Fragment

The Internet Protocol (IP) implements datagram "fragmentation", so that packets may be formed that can pass through a link with a smaller MTU (maximum transmission unit) than the original datagram size. IETF RFC 791 describes a procedure for IP fragmentation, and transmission and reassembly of datagrams. RFC 815 describes a similar reassembly algorithm. The details of the fragmentation mechanism, as well as the overall architectural approach to fragmentation, are different between IPv4 and IPv6. IPv4 hosts must make a best-effort attempt to reassemble fragmented IP datagrams with a total reassembled size of up to 576 bytes (equal to the minimum MTU for IPv4). IPv4 hosts may also attempt to reassemble fragmented IP datagrams larger than 576 bytes, but they are also permitted to silently discard such larger datagrams.

(In IPv6, this minimum capability is increased to 1280 bytes - larger than the minimum MTU for IPv4.)

This is the IPv4 Fragment for matching the ECE. This involves the settings for the More Fragments (MF) bit and the Fragment Offset (FRAG OFFSET) field for an IPv4 frame.

The valid values are:

Any: The ECE will match any MF bit.

Fragment: IPv4 frames where the MF bit is set or the FRAG OFFSET field is greater than zero must be able to match this entry. Fragmentation of an internet datagram is necessary when it originates in a local net that allows a large packet size and must traverse a local net that limits packets to a smaller size to reach its destination.

Non-Fragment: 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. An internet datagram can be marked "don't fragment." Any internet datagram so marked is not to be internet fragmented under any circumstances. If an internet datagram marked 'don't fragment' cannot be delivered to its destination without fragmenting it, it is to be discarded instead.

Protocol Value

Select the IPv4 protocol value to be used. The default is **0**. The valid range is **0-255**. This field displays only if "Other" is selected in the IPv4 "Protocol" field above.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Cancel: Return to the previous page; any changes made locally will be undone.

Click the **Save** button when done.

Verify your ECE Control List configuration.

ECE ID	UNI Ports	Ingress Matching					Actions				Conflict	
		Tag Type	VID	PCP	DEI	Frame Type	Direction	EVC ID	Tag Pop Count	Policy ID		Ingress Map
1	3,4	Any	Any	Any	Any	Any	Both	1	0	0	Disabled	No
2	9,10	Untagged	-	-	-	Any	Both	2	0	0	Disabled	No
3	9,10	Any	Any	Any	Any	Any	Both	2	0	0	Disabled	No

You can use the Modification buttons at the far right-hand side of the table to modify the table.

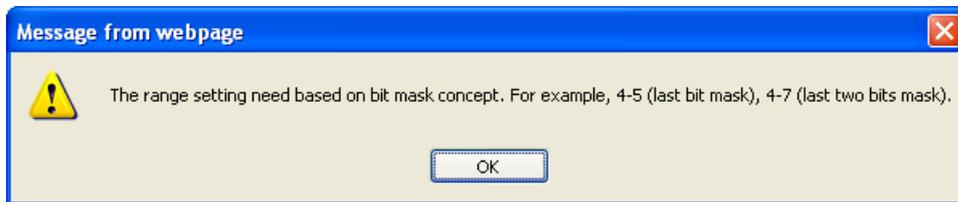
Modification Buttons

You can modify each ECE (EVC Control Entry) in the table using the following buttons:

- : Inserts a new ECE before the current row.
- : Edits the ECE row.
- : Moves the ECE up the list.
- : Moves the ECE down the list.
- : Displays the webpage message "Do you want to delete this entry?" Click **OK** to delete the ECE.
- : The lowest plus sign adds a new entry at the bottom of the ECE listings.

Messages

Message: The range setting need based on bit mask concept. For example, 4-5 (last bit mask), 4-7 (last two bits mask).



Meaning: At the “**Inner Tag VLAN ID Range**” field in the “**UNI Matching**” section you entered an invalid value and the clicked the **Save** button.

Recovery:

Enter the range setting need based on bit mask concept. Enter 4-5 (last bit mask) or 4-7 (last two bits mask).

Click the **Save** button.

Continue operation.

L2CP (Layer 2 Control Protocol Processing)

L2CP Processing for EPL per MEF 6.1

The MEF introduced the L2CP attribute for its Ethernet services and UNI in MEF 6. Each of the Layer 2 control protocols can be tunnelled, forwarded, or discarded.

Note: You can use the “`evc port l2cp`” command, but not the LIB-44xx web interface, to set or show port L2CP mode the L2CP settings. See the LIB-44xx CLI Reference manual for specifics.

The L2CP processing for EPL per MEF 6.1 is outlined below:

Protocol	MAC DA	L2CP Requirement		Applicability
		Option 1	Option 2	
STP/RSTP/MSTP	01-80-C2-00-00-00	Must Tunnel	Must Tunnel	All UNIs in the EVC
PAUSE	01-80-C2-00-00-01	Should Discard	Should Discard	All UNIs in the EVC
LACP/LAMP	01-80-C2-00-00-02	Should Peer or Discard	Should Tunnel	Option 1: Per UNI Option 2: All UNIs in the EVC
Link OAM	01-80-C2-00-00-02	Should Peer or Discard	Should Tunnel	Option 1: Per UNI Option 2: All UNIs in the EVC
Port Authentication	01-80-C2-00-00-03	Should Peer or Discard	Should Tunnel	Option 1: Per UNI Option 2: All UNIs in the EVC
E-LMI	01-80-C2-00-00-07	Should Peer or Discard	Must Tunnel	Option 1: Per UNI Option 2: All UNIs in the EVC
LLDP	01-80-C2-00-00-0E	Should Discard	Must Tunnel	All UNIs in the EVC
GARP Block	01-80-C2-00-00-20 through 01-80-C2-00-00-2F	Must Discard or Tunnel	Must Tunnel	All UNIs in the EVC

The L2CP processing for EVPL service per MEF 6.1 is outlined below:

Protocol	MAC DA	L2CP Requirement	Applicability
STP/RSTP/MSTP	01-80-C2-00-00-00	Must Peer or Discard	All UNIs in the EVC
PAUSE	01-80-C2-00-00-01	Must Discard	All UNIs in the EVC
LACP/LAMP	01-80-C2-00-00-02	Must Peer or Discard	Per UNI
Link OAM	01-80-C2-00-00-02	Must Peer or Discard	Per UNI
Port Authentication	01-80-C2-00-00-03	Must Peer or Discard	Per UNI
E-LMI	01-80-C2-00-00-07	Must Peer or Discard	Per UNI
LLDP	01-80-C2-00-00-0E	Must Discard	All UNIs in the EVC
GARP	01-80-C2-00-00-20 through 01-80-C2-00-00-2F	Must Peer, Tunnel or Discard	Per UNI

At LIB-44xx v 1.0.2, L2CP disposition is available via CLI on a per-port basis. Upcoming releases will provide L2CP at the per-EVC level.

The LIB-44xx L2CP settings do not map exactly as the MEF defines; the table below shows the LIB-44xx L2CP behaviour. There is no explicit way to ‘discard’, but only to redirect to CPU and if no

protocol is instantiated on that port, then it is discarded. The BPDU/GARP behaviour works a little different under normal mode as well (more to the IEEE recommendation).

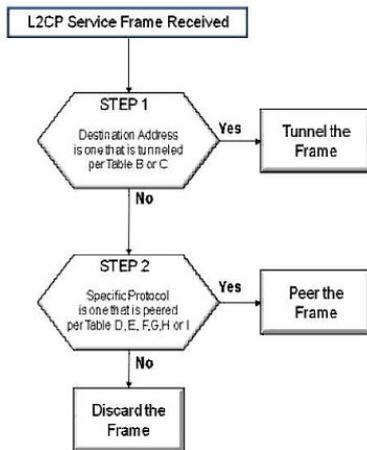
The L2CP MIB is based on the MEF-UNI MIB. Use the table below to decode the L2CP settings.

L2CP	Normal	Redirect	Forward
BPDU	Discards (or Peers, depending on if the protocol is enabled on that port)	Peers (or discards depending on if the protocol is enabled on that port)	Tunnel
GARP	Tunnel	Discards (since GARP protocols are not supported)	Tunnel

At LIB-44xx software version 1.2.3, the L2CP default changed from 'Normal' to (BPDU -> Peer, GARP -> tunnel) for both the LIB-44xx CLI and SNMP.

L2CP per MEF 6.1.1

MEF 6.1.1 expands on the EVC L2CP handling rules with L2CP Processing Requirements for Service Frames with a MAC DA in the range of 01-80-C2-00-00-00 to -0F. The action ('tunnel', 'peer', or 'discard') for each L2CP Service Frame is decided using a two-step logic based on the frame's MAC DA, and then its Ethertype and subtype or LLC code (i.e., based on the protocol). The logic for processing of the L2CP Service Frames is shown below:



See http://metroethernetforum.org/Assets/Technical_Specifications/PDF/MEF_6.1.1.pdf for more information.

OVC Layer 2 Control Protocol Tunnelling per MEF 26.0.2

The MEF introduced L2CP Tunnelling per MEF 26.0.2 as the process by which a frame containing a Layer 2 Control Protocol is transferred between External Interfaces.

Layer 2 Control Protocol Tunnelling

The Layer 2 Control Protocol Service Frame is described in MEF 10.2. An ENNI Frame with a Destination MAC Address (Table B) is defined to be a Layer 2 Control Protocol ENNI Frame. Other ways of denoting a Layer 2 Control Protocol ENNI Frame at a given ENNI can be agreed to by the two Operators involved in the given ENNI.

Table B - MAC Addresses that Identify a Layer 2 Control Protocol ENNI Frame

MAC Addresses
01-80-C2-00-00-00 through 01-80-C2-00-00-0F
01-80-C2-00-00-20 through 01-80-C2-00-00-2F
01-80-C2-00-00-10

Table C – Format Relationships for Tunnelled L2CP Service and ENNI Frames

Ingress Interface	Egress Interface	Egress Frame Format*
UNI (L2CP Service Frame)	UNI (L2CP Service Frame)	Identical to the ingress frame.
UNI (L2CP Service Frame)	ENNI (L2CP ENNI Frame)	All fields from the Destination Address through the Payload of the ingress Service Frame present and unchanged. S-Tag added in after the Source Address.
ENNI (L2CP ENNI Frame)	UNI (L2CP Service Frame)	All fields from the Destination Address through the Payload except the S-Tag of the ingress ENNI Frame present and unchanged. No S-Tag is present.
ENNI (L2CP ENNI Frame)	ENNI (L2CP ENNI Frame)	All fields from the Destination Address through the Payload of the ingress ENNI Frame present. The content of the S-Tag can be changed while all other fields are unchanged.

[1] When an L2CP Service Frame or L2CP ENNI Frame is tunnelled, the frame **MUST** be delivered to all OVC End Points, other than the ingress OVC End Point, that are associated by the OVC and the format relationships detailed in Table C **MUST** be maintained.

[2] An ingress L2CP Service Frame that is not mapped to an existing OVC End Point **MUST NOT** be tunnelled.

[3] An ingress L2CP ENNI Frame that is not mapped to an existing OVC End Point **MUST NOT** be tunnelled. This means that an ingress L2CP ENNI Frame that does not have an S-Tag is not to be tunnelled because the Operator has no information on which OVC to use to tunnel the frame.

* The MEF further notes that “The Frame Check Sequence in the egress frame might need to be recalculated.”

For More Information

See the MEF Technical Specifications at http://metroethernetforum.org/page_loader.php?p_id=29.

Performance Monitoring Configuration

Performance Monitor will continuously check the bandwidth usage of each instance, if a threshold crossing alarm condition occurs, the system will generate TCA traps to server while SNMP trap is enabled. Up to 132 Instances are supported.

Performance Monitoring is done via collection of Performance information in Measurement Intervals. The information can be store in non-volatile memory and can be transferred to a server. Bandwidth thresholds are used to monitor a certain traffic usage and traps would be generated to EMS server if threshold crossing condition occurs. Five PM Sessions are currently supported: Loss Measurement (LM), Delay Measurement (DM), EVC, and ECE. The Performance Monitor (PM) will continuously check the bandwidth usage of each instance; if a threshold crossing alarm condition occurs, the system will generate TCA traps to server while SNMP trap is enabled. Up to 132 Instances are supported.

The PM Session gathers all PM data sets for all enabled PM features in Measurement Intervals along with other system related information. The limit of Measurement Intervals before overwrite is 96 per PM feature. When the Measurement Interval limit is reached, the oldest Measurement Interval is overwritten. Up to 64 data sets is supported for each Measurement Interval. A Measurement Interval is configurable for each PM session, in the range of 1 - 60 minutes; the default is 15 minutes.

The overall PM configuration process includes configuring to:

Collect stats at the **Configuration > Performance Monitor > Configuration** menu path.

Push the gathered PM Data Sets for all Measurement Intervals to a TFTP server at **Configuration > Performance Monitor > Transfer Mode**.

Inspect the PM statistics from **Monitor > Performance Monitor** menu path when threshold is reached and a trap is generated. Select LM, DM, EVC, or ECE Statistics to view.

View the PM Measurement Interval Information for a selected Information Type (LM, DM, EVC, ECE, or Port) at **Monitor > Performance Monitor > Interval Information**.

> Configuration > Performance Monitor > Configuration

The **Configuration > Performance Monitor > Configuration** menu path lets you enable PM session and storage.

Type	Enable Session	Enable Storage	Measurement Interval(mins)
Loss Measurement	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	15
Delay Measurement	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	15
Delay Measurement Binning	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	
EVC	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	15

The PM Session and Storage Configuration parameters are described below:

Type

The type of performance monitor data (Loss Measurement, Delay Measurement, EVC, or ECE).

Enable Session

Enable or disable the performance monitor session.

Enable Storage

Enable or disable the performance monitor storage.

Measurement Interval(mins)

The measurement interval for the performance monitor in minutes. The valid range is 1-60. The default is 15.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Performance Monitor Transfer Mode Configuration

The **Configuration > Performance Monitor > Transfer Mode** menu path lets you perform PM Transfer Configuration.

All gathered PM Data Sets for all Measurement Intervals can be pushed to a TFTP server if that server is enabled. This can be done at configurable point in time (the configured schedule hours combined with schedule minutes). The transfer of PM data must be done in a separate file per PM feature (LM, DM, EVC Statistics, ECE Statistics and Port Statistics). All transferred files are zipped using [GZIP](#). The naming of the transferred files containing 'All' Measurement intervals is:

- 'System Name'_history_'Year'-'day'-month'_'hour'h'min.'m'sec.'s_DM_'all'.cvs.gz
- 'System Name'_history_'Year'-'day'-month'_'hour'h'min.'m'sec.'s_LM_'all'.cvs.gz
- 'System Name'_history_'Year'-'day'-month'_'hour'h'min.'m'sec.'s_EVC_'all'.cvs.gz
- 'System Name'_history_'Year'-'day'-month'_'hour'h'min.'m'sec.'s_ECE_'all'.cvs.gz
- 'System Name'_history_'Year'-'day'-month'_'hour'h'min.'m'sec.'s_PORT_'all'.cvs.gz

The format of the transferred files is CVS (Comma Separated Values). The first line in the file must list the names of the parameters in the following lines. Each of the following lines lists the values of the PM data sets contained in this file.

The parameters are described below:

PM Transfer Mode

Configure the operation mode per system. Possible modes are:

Enabled: Enable the PM Transfer Mode function.

Disabled: Disable PM Transfer Mode (default mode).

Scheduled Hours

Here you can select one **or more** of the 24 hours in a day, when PM data transfer will happen.

00:00

01:00

02:00

03:00

04:00

05:00

06:00

07:00

:::::

20:00

21:00

22:00

23:00

The default is none selected. Use the keyboard Shift key or Ctrl key and mouse click to select multiple Schedule hours.

Scheduled Minutes

Here you can select one **or more** of the four 15 minute quarters of an hour when the PM data transfer will occur.

00:00

00:15

00:30

00:45

The default is none selected.

Scheduled Offset

Here you can configure a fixed offset that is added to the scheduled transfer time.

The range is **0-15** minutes. The default is **0** minutes.

The sum of the Scheduled Fixed Offset plus the Scheduled Random Offset must not exceed **15** minutes.

Random Offset

It must be possible to configure a random offset that is added to the scheduled transfer time.

The offset added to the scheduled transfer time must be a random value in the range 0-Scheduled Offset.

The valid range is **0-900** seconds. The default is **0** seconds.

The sum of the Scheduled Offset plus the Random Offset must not exceed **15** minutes.

Server Directory URL

Enter the full URL of the TFTP server and the corresponding directory (if any) for uploading.

Enable the TFTP server by entering `tftp://` followed by the domain name or IP address of the TFTP server (e.g., `tftp://ip-addr` or `tftp://dns-name`).

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Example

Make sure that the TFTP server is configured and running. The LIB-4400 Server Directory URL must match the TFTP Server Interface setting.

The screenshot displays the NET2EDGE web interface for a LIB-4424 Carrier Ethernet Network Interface Device. The main configuration area is titled "PM Transfer Configuration" and includes the following settings:

- PM Transfer Mode:** Enabled
- Scheduled hours:** 00:00 HRS, 01:00 HRS, 02:00 HRS, 03:00 HRS, 04:00 HRS, 05:00 HRS, 06:00 HRS
- Scheduled minutes:**
 - Every 00:00
 - Every 00:15
 - Every 00:30
 - Every 00:45
- Scheduled offset:** 0 minutes
- Random offset:** 0 seconds
- Server Directory URL:** `tftp://192.168.1.185`
- Transfer Interval Mode:**
 - All available intervals
 - New intervals since last
 - Fixed number of intervals
- Transfer Option:**
 - Include intervals from

At the bottom of the configuration area are "Save" and "Reset" buttons. An inset window titled "SolarWinds TFTP Server" is shown, displaying the status "TFTP Server Service Started [UDP - 69]".

QoS Configuration

LIB-44xx QoS Configuration is performed from the **Configuration > QOS** menu path. 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 secret to a successful end-to-end business solution. QoS provides a set of techniques to manage network resources to achieve this success.

Bandwidth Profiling

MEF 10.2 defines 'a bandwidth profile' as "a method of characterizing Service Frames for the purpose of rate enforcement or policing." A sample bandwidth profile use case for a provider dropping a port at an enterprise customer for triple-play services (video, voice and data). The provider may have a Service level agreement with a customer for 5M for video, 3M for voice and 4M for data service, but the overall service cannot exceed 7M for that subscriber.

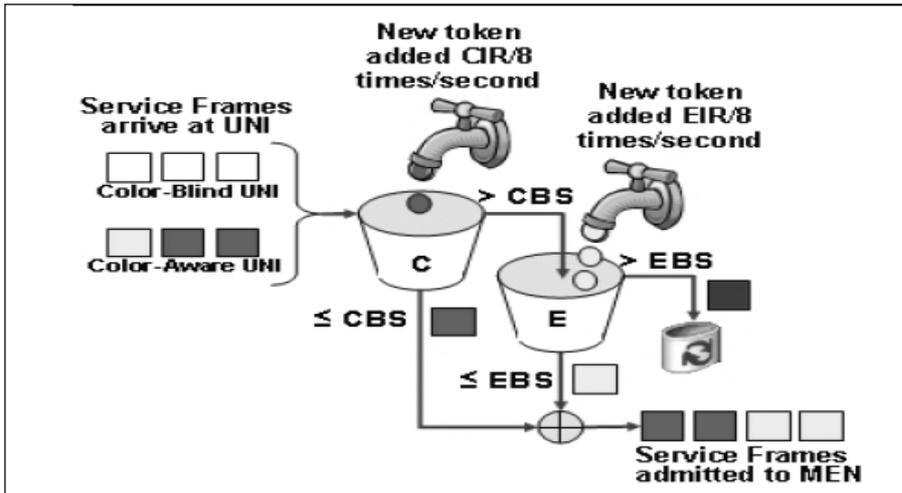


Figure 20. Colour Aware Dual leaky bucket for Bandwidth Profiling

The LIB-44xx device supports bandwidth profiling per MEF 10.2, section 7.11 at three levels:

- Ingress bandwidth profile per UNI (port),
- Ingress bandwidth profile per EVC (VLAN) per UNI, and
- Ingress bandwidth profile per Cos per EVC per UNI

For options 2 and 3, the LIB-44xx can provide an overall UNI bandwidth profile as well.

The figure below illustrates a provider’s SLA with a customer for 5M video / 3M voice / 4M data service where the overall service cannot exceed 7M for that subscriber.

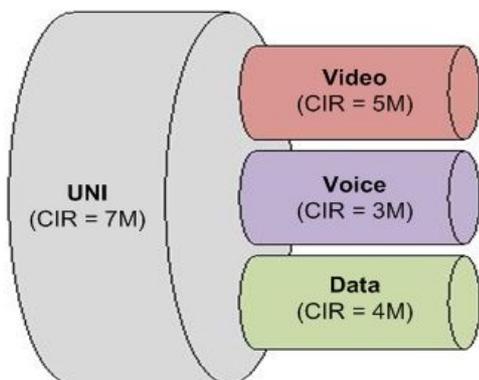
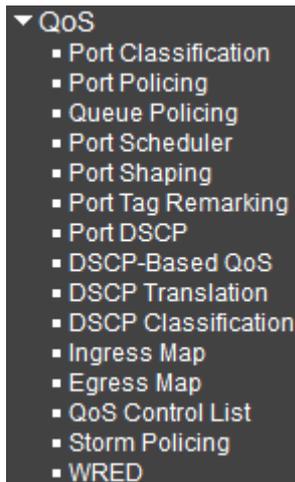


Figure 21. Example SLA for Bandwidth profiling



QoS Configuration Sub-menus

LIB-44xx QoS configuration sub-menus include Port Classification, Port Policing, Queue Policing, Port Scheduler, Port Shaping, Port Tag Remark, Port DSCP, DSCP-Based QoS, DSCP Translation, DSCP Classification, QoS Control List, Storm Control, and WRED configuration.

Note that the functions configured at **Configuration > QoS** are monitored at the **Monitor > Ports** menu path. For example:

Configure at

Configuration > QoS > QoS Control List
 Configuration > QoS > Port Policing
 Configuration > QoS > Queue Policing

Monitor at

Monitor > Ports > QCL Status
 Monitor > Ports > Detailed Statistics
 Monitor > Ports > QoS Statistics

Each of these QoS configuration sub-menus is explained below:

Port Classification

The **Configuration > QoS > Port Classification** menu path displays the QoS Ingress Port Classification table. This page lets you configure the basic QoS Ingress Classification settings for all LIB-44xx ports.

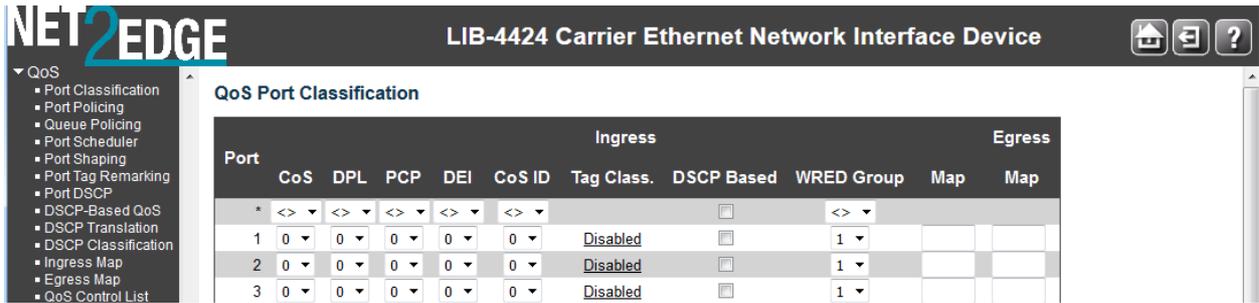
All LIB-44xx frames ingressing the device are associated with a priority to help with classification into the output queues on the egress port and also when to transmit the frame by the egress scheduler.

The LIB-44xx provides a set of basic classification and advanced classification using the TCAM.

This section lists the various classifications, including:

1. Ingress Port Priority,
2. Priority Code Point and Drop Eligible Indicator bits (IEEE Priority),
3. DSCP Traffic class (IP priority),
4. DSCP translation, and
5. Advanced per flow QoS.

1. Ingress Port Priority: For frames which don't have any priority fields (layer 2 or layer 3), by default take the priority of the ingress port. Each port is associated with a default priority fields which is user configurable.



The displayed Port Classification settings are explained below:

Port

The port number for which the configuration below applies (e.g., ports 1-4 on the LIB-4400). The * in the Port column acts as a ‘wild card’ character which causes the selections in this row to be applied to all other Ports (rows) in the table for which this selection is valid. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

QoS class

Controls the default QoS class (the QoS class for frames not classified in any other way). There is a one to one mapping between QoS class, queue and priority. A QoS class of 0 (zero) has the lowest priority.

Select 0 - 7 at the dropdown.

If the port is VLAN aware, the frame is tagged and Tag Class. is enabled, then the frame is classified to a QoS class that is mapped from the PCP and DEI value in the tag. Otherwise the frame is classified to the default QoS class.

The classified QoS class can be overruled by a QCL entry.

Note: If the default QoS class has been dynamically changed, then the actual default QoS class is shown in parentheses after the configured default QoS class.

DP level

Controls the default Drop Precedence Level.

DPL = 0 (zero) corresponds to 'Committed' (Green) frames. Packets with a DPL of 0 are least likely to be dropped

DPL = 1 or higher corresponds to 'Discard Eligible' (Yellow) frames.

DPL = 2 packets with a DPL of 2 are second most likely to be dropped.

DPL = 3 packets with a DPL of 3 are most likely to be dropped.

If congestion occurs within a class, the packets with the higher drop precedence are discarded first. To prevent issues associated with tail drop, more sophisticated drop selection algorithms such as random early detection (RED) can be used.

Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

All frames are classified to a DP level.

1) If the Port is VLAN aware, if the frame is tagged and then if it is not tagged, then if:

Tagged: From tag classification mapping for port if enabled;

Untagged: Use default QoS class and DP level for the port.

2) If the Port is VLAN unaware, if the frame is tagged and then if it is not tagged, then use default QoS class and DP level for port.

The classified DP level can be overruled by a QCL entry.

PCP

Controls the default PCP value. All frames are classified to a 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.

Select **0** - **7** at the dropdown.

DEI

Controls the default DEI value. DEI (Drop Eligible Indicator) is a 1-bit field in the VLAN tag. All frames are classified to a 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.

Select **0** or **1** at the dropdown.

Tag Class.

Shows the classification mode for tagged frames on this port.

Disabled: Use default QoS class and DP level for tagged frames.

Enabled: Use mapped versions of PCP and DEI for tagged frames.

Click on the mode (Enabled or Disabled) in order to configure the mode and/or mapping (described below).

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 QoS class and DP level.

DSCP Based

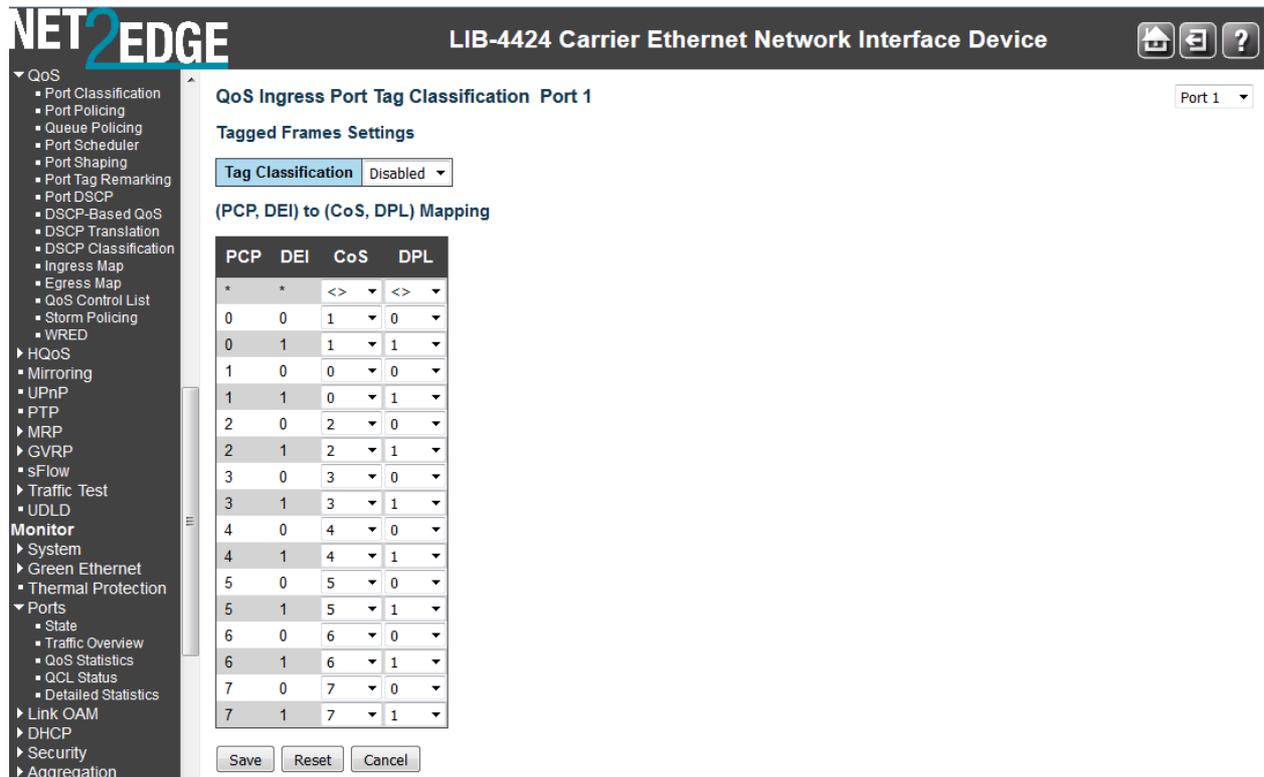
Check the checkbox to enable DSCP Based QoS Ingress Port Classification. DSCP (Differentiated Services Code Point) is a field in the header of IP packets for packet classification purposes. The default is unchecked (disabled).

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

When you click the **Disabled** link in the **Tag Class.** column in a row for a port, the Tag Classification page displays for that port.



Here you can:

- Set **Tag Classification** to Enabled (the default is Disabled),
- Set the (PCP, DEI) to (QoS class, DP level) Mapping:
- Set QoS class** to 0-7, and/or
- Set DP level** to 0 or 1.

You can click the browser’s Back button to go back to the QoS Ingress Port Classification page. Click the **Save** button when done; the QoS Ingress Port Classification page displays again with the new settings.

2. Priority Code Point and Drop Eligible Indicator bits (IEEE Priority): The PCP and DEI bits in tagged frame can be used for QoS and a table of PCP, DEI to QoS class is programmable per port. Ingress Policers and Shapers

The LIB-44xx supports up to 256 Policers that are programmable to any port and any flow. Each of these policers is compliant with MEF dual leaky buckets and they can be made colour-blind or colour aware. The individual statistics of each policer can be obtained. Each frame can be policed by up to 3 policers based on port, queue (based on QoS) and any ACE (access control entry) rules. Each port is also equipped with an ingress shaper which controls the rate of transfer between ingress and egress port queues.

Port Policing

The **Configuration > QoS > Port Policing** menu path displays the QoS Ingress Port Policers table. This page allows you to configure the Policer settings for all LIB-44xx ports. The policer can limit the bandwidth of received frames. It is located in front of the Ingress queue.

Port	Enable	Rate	Unit	Flow Control
*	<input type="checkbox"/>	500	<>	<input type="checkbox"/>
1	<input checked="" type="checkbox"/>	500	kbps	<input type="checkbox"/>
2	<input type="checkbox"/>	500	kbps	<input type="checkbox"/>
3	<input type="checkbox"/>	500	kbps	<input type="checkbox"/>
4	<input type="checkbox"/>	500	kbps	<input type="checkbox"/>
5	<input type="checkbox"/>	500	kbps	<input checked="" type="checkbox"/>
6	<input type="checkbox"/>	500	kbps	<input checked="" type="checkbox"/>

The displayed settings are explained below:

Port

The port number for which the configuration below applies (e.g., ports 1-4 on the LIB-4400). The * in the Port column acts as a 'wild card' character which causes the selections in this row to be applied to all other Ports (rows) in the table for which this selection is valid. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Enabled

Controls whether the policer is enabled on this LIB-44xx port. Check the checkbox to enable port policing on this port (row). The default is unchecked (disabled).

Rate

Controls the rate for the policer. The default value is "500". This value is restricted to 100-1000000 when the "Unit" is "kbps" or "fps", and it is restricted to 1-13200 when the "Unit" is "Mbps" or "kfps".

Unit

Controls the unit of measure for the policer rate as **kbps**, **Mbps**, **fps** or **kfps**. The default value is "**kbps**". (Where 'bps' is bits per second, and 'fps' is frames per second.)

Flow Control

Check or uncheck the checkbox to enable or disable Flow Control on a per-port basis. If flow control is enabled and the port is in flow control mode, then pause frames are sent instead of discarding frames.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Egress Shaping

The LIB-44xx has port-level and queue-level shapers. The shapers use the ‘leaky bucket’ method of rate limiting, and have a fixed size buffer to allow bursts over which frames will be discarded. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

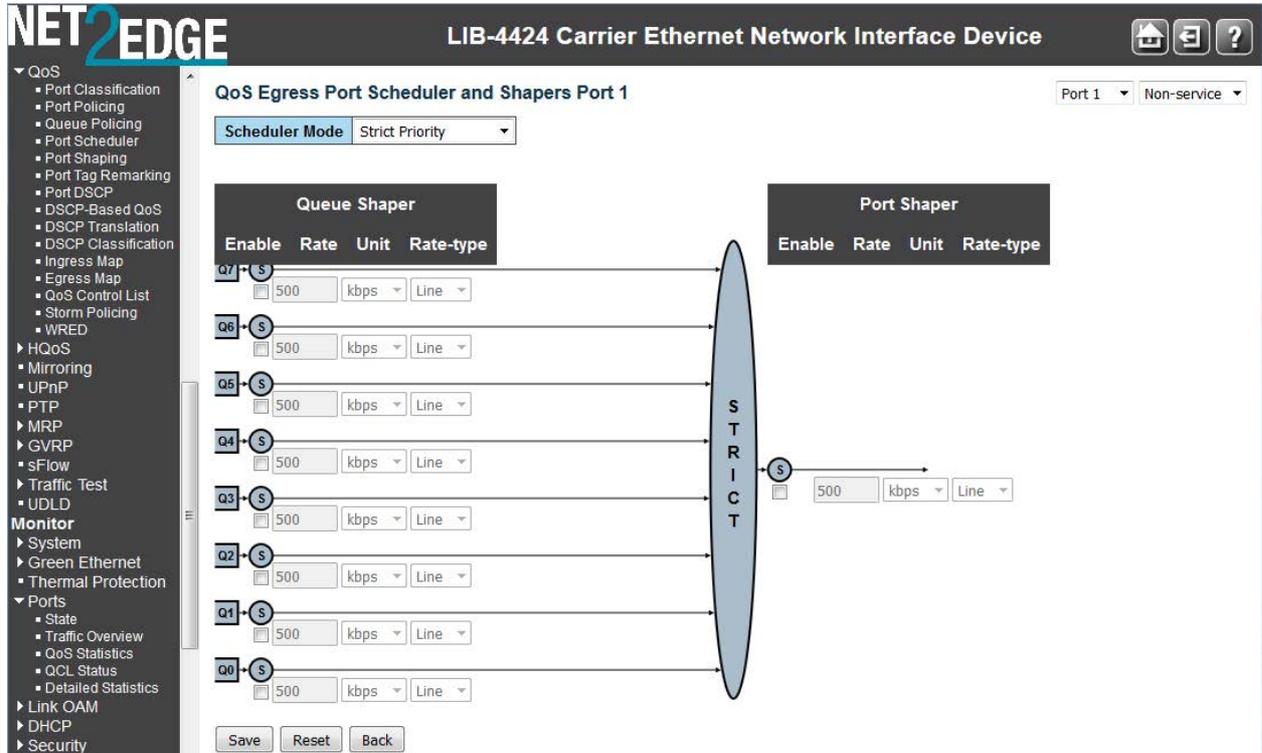
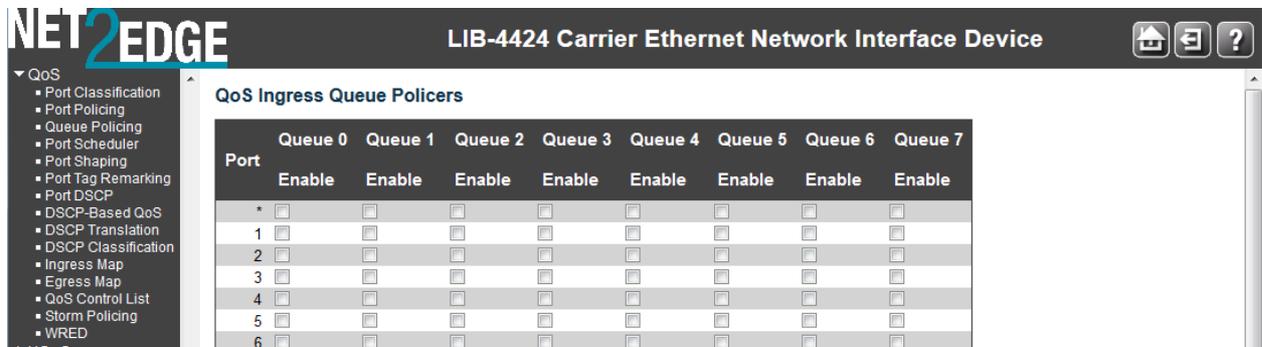


Figure 22. Egress Shaper on Port 1

Queue Policing

The **Configuration > QoS > Queue Policing** menu path lets you configure the Queue Policer settings for all LIB-44xx ports.

Each policer can limit the bandwidth of received frames. A policer is located in front of the ingress queue.



Note that the functions configured at **Configuration > QoS** are monitored at the **Monitor > Ports** menu path. For example:

Configure at

Monitor at

Configuration > QoS > QoS Control List Monitor > Ports > QCL Status

Configuration > QoS > Port Policing

Monitor > Ports > Detailed Statistics

Configuration > QoS > Queue Policing

Monitor > Ports > QoS Statistics

The QoS Ingress Queue Policers parameters are explained below:

Port

The port number for which the configuration below applies (e.g., ports 1-4 on the LIB-4400). The first row in the table, marked by the * sign, enables or disables (checks or unchecks the ‘enabled’ checkbox) for all of the table rows. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Queue x Enable

Check the checkbox to enable the queue policer on this LIB-44xx port (e.g., 1-4). This expands the table to let you define the rate(s) as shown below:

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

The screen below shows multiple Queue Policers configured for LIB-44xx port 2. The configuration below has Port 2 with Queues 0, 1, and 2 enabled and set for a rate of 1 Mbps. Note that queues 3-7 are not enabled.

Port	Queue 0			Queue 1			Queue 2			Queue 3	Queue 4	Queue 5	Queue 6	Queue 7
	E	Rate	Unit	E	Rate	Unit	E	Rate	Unit	Enable	Enable	Enable	Enable	Enable
*	<input type="checkbox"/>	500	<>	<input type="checkbox"/>	500	<>	<input type="checkbox"/>	500	<>	<input type="checkbox"/>				
1	<input type="checkbox"/>	500	kbps	<input type="checkbox"/>	500	kbps	<input type="checkbox"/>	500	kbps	<input type="checkbox"/>				
2	<input checked="" type="checkbox"/>	500	kbps	<input checked="" type="checkbox"/>	500	kbps	<input checked="" type="checkbox"/>	500	kbps	<input type="checkbox"/>				
3	<input checked="" type="checkbox"/>	500	kbps	<input checked="" type="checkbox"/>	500	kbps	<input checked="" type="checkbox"/>	500	kbps	<input type="checkbox"/>				
4	<input type="checkbox"/>	500	kbps	<input type="checkbox"/>	500	kbps	<input type="checkbox"/>	500	kbps	<input type="checkbox"/>				
5	<input type="checkbox"/>	500	kbps	<input type="checkbox"/>	500	kbps	<input type="checkbox"/>	500	kbps	<input type="checkbox"/>				
6	<input type="checkbox"/>	500	kbps	<input type="checkbox"/>	500	kbps	<input type="checkbox"/>	500	kbps	<input type="checkbox"/>				

The QoS Ingress Queue Policers parameters are explained below:

Port

The port number for which the configuration below applies. The first row in the table, marked by the * sign, enables or disables (checks or unchecks the ‘enabled’ checkbox) for all of the table rows.

Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

E (Enabled)

Controls whether the queue policer is enabled (checkbox checked) on this LIB-44xx port. The E column only displays if the checkbox is checked.

Rate

Controls the rate for the queue policer. The default value is **500**. This value is restricted to **100 - 1000000** when the "Unit" is "kbps", and it is restricted to **1 - 3300** when the "Unit" is "Mbps". This field only displays if one or more of the queue policers are enabled.

Unit

Controls the unit of measure for the queue policer rate as **kbps** or **Mbps**. The default value is "**kbps**". This field only displays if one or more of the queue policers are enabled.

Qx Enable

Controls whether this queue (Queue 1 - Queue 7) is enabled (checkbox checked) on this LIB-44xx port.

Port Scheduler

The **Configuration > QoS > Port Scheduler** menu path displays the QoS Egress Port Schedulers table.

Egress Scheduler and Shaper: Each port has an egress scheduler and a set of egress shapers. The scheduler on each port can operate in strict priority (the default) or in a mixed mode where the high priority queues are Strict and the rest are in DWRR (Deficit Weighted Round Robin). The scheduler mode in DWRR lets you assign weights to the individual QoS queues. The Egress shapers available on a per-QoS queue basis provide shaping at a least granular level, and a port egress shaper provides an overall shaping and rate limiting the throughput.

DWRR (Deficit Weighted Round Robin)

The DWRR uses a cost-based algorithm (compared to a weight-based algorithm). A high cost implies a small share of the bandwidth. When DWRR is enabled, each of the queues 5 through 0) are programmed with a cost (a number from 1 - 32). The programmable DWRR costs determine the behaviour of the DWRR algorithm. The costs result in weights for each queue. The weights are relative to one another, and the resulting share of the egress bandwidth for a particular QoS class is equal to the queue's weight divided by the sum of all the queues' weights.

Costs can be converted to weights (and vice versa) with these two algorithms:

Weight to Cost: given a desired set of weights ($W_0, W_1, W_2, W_3, W_4, W_5$), calculate the costs using the following algorithm:

1. Set the cost of the queue with the smallest weight ($W_{smallest}$) to cost 32.
2. For any other queue Q_n with weight W_n , set the corresponding cost C_n to:

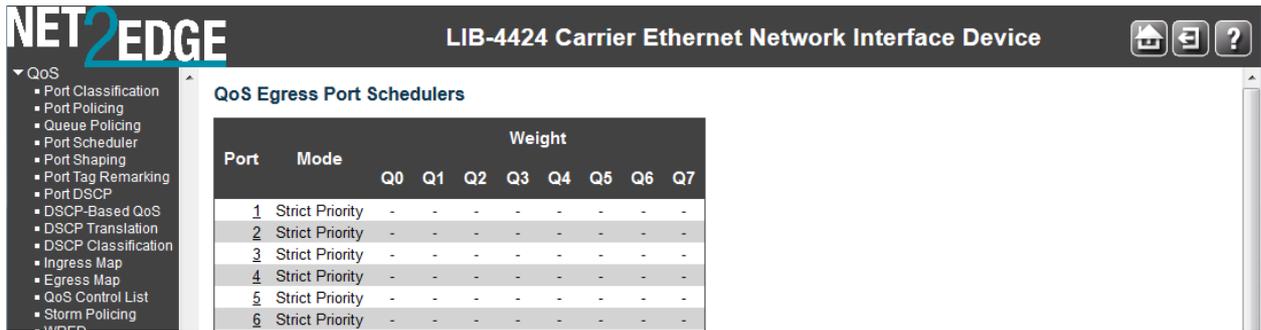
$$C_n = 32 \times W_{smallest} / W_n$$

Cost to Weight: given a set of costs for all queues ($C_0, C_1, C_2, C_3, C_4, C_5$), the calculate the resulting weights using the following algorithm:

1. Set the weight of the queue with the highest cost ($C_{highest}$) to 1.
2. For any other queue Q_n with cost C_n , set the corresponding weight W_n to:

$$W_n = C_{highest} / C_n$$

The **Configuration > QoS > Port Scheduler** menu path provides an overview of QoS Egress Port Schedulers for all LIB-44xx ports.



The displayed QoS Egress Port Schedulers settings are described below:

Port

The logical port for the settings contained in the same row (e.g., ports 1-4 on the LIB-4400). Click on the port number in order to configure the schedulers. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Mode

Shows the scheduling mode for this port ("Strict Priority" or "Weighted"). Generally, Strict Priority (SP) queues are scheduled before WRR queues.

Strict Priority - The port transmits all packets out of higher priority queues before transmitting any from the lower priority queues.

Weighted - (WRR or Weighted Round Robin) - The port transmits a set number of packets from each queue, in a round robin fashion, so that each has a chance to transmit traffic.

Weight (Q0 - Q5)

Shows the weight for this queue (Q0-Q5) and port (port 1-4).

WRR setting examples for the number of packets transmitted from each queue are shown below:

These values are permanent and you cannot change them.

Port Egress Queue Max. No. of Packets

Q3	8
Q2	4
Q1	2
Q0	1

When you click on the Port number in a row, the QoS Egress Port Scheduler and Shapers for that Port display (e.g., for Port 1 in the screen sample below).

NET2EDGE LIB-4424 Carrier Ethernet Network Interface Device

QoS Egress Port Scheduler and Shapers Port 1

Scheduler Mode: Strict Priority

Queue Shaper			
Enable	Rate	Unit	Rate-type
<input checked="" type="checkbox"/>	500	kbps	Line
<input checked="" type="checkbox"/>	500	kbps	Line
<input checked="" type="checkbox"/>	500	kbps	Line
<input checked="" type="checkbox"/>	500	kbps	Line
<input checked="" type="checkbox"/>	500	kbps	Line
<input checked="" type="checkbox"/>	500	kbps	Line
<input checked="" type="checkbox"/>	500	kbps	Line
<input checked="" type="checkbox"/>	500	kbps	Line

Port Shaper			
Enable	Rate	Unit	Rate-type
<input checked="" type="checkbox"/>	500	kbps	Line

Diagram: Queue Shaper (Q0-Q7) → STRICT Scheduler → Port Shaper

Buttons: Save, Reset, Back

If **Scheduler Mode** is set to “Weighted”, then the Port Scheduler and Shapers for that Port display (e.g., for Port 1 as shown below).

NET2EDGE LIB-4424 Carrier Ethernet Network Interface Device

QoS Egress Port Scheduler and Shapers Port 1

Scheduler Mode: 5 Queues Weighted

Queue Shaper			
Enable	Rate	Unit	Rate-type
<input checked="" type="checkbox"/>	500	kbps	Line
<input checked="" type="checkbox"/>	500	kbps	Line
<input checked="" type="checkbox"/>	500	kbps	Line
<input checked="" type="checkbox"/>	500	kbps	Line
<input checked="" type="checkbox"/>	500	kbps	Line
<input checked="" type="checkbox"/>	500	kbps	Line
<input checked="" type="checkbox"/>	500	kbps	Line
<input checked="" type="checkbox"/>	500	kbps	Line

Queue Scheduler	
Weight	Percent
17	20%
17	20%
17	20%
17	20%
17	20%
17	20%
17	20%
17	20%

Port Shaper			
Enable	Rate	Unit	Rate-type
<input checked="" type="checkbox"/>	500	kbps	Line

Diagram: Queue Shaper (Q0-Q7) → D W R R Scheduler → STRICT Scheduler → Port Shaper

Buttons: Save, Reset, Back

The Port Scheduler and Shapers for a specific port are configured on this page. The parameters are explained below:

Scheduler Mode

Controls whether the scheduler mode is "**Strict Priority**" or "**Weighted**" on this LIB-44xx port ("Weighted" selection shown above).

Strict Priority : Select Strict Priority to process the packets with the highest priority first (strictly according to priority).

Weighted : Select WRR (Weighted Round-Robin) to process packets according to the weight of each priority. When a priority level has reached its egress weight, the system will process the packets in the next level even if there are remaining packets.

Queue Shaper Enable

Controls whether the queue shaper is enabled for this queue on this LIB-44xx port. Check the checkbox to turn on the Queue Shaper function for this port. This sets the rate limit for this port (the actual rate set will depend on the "Unit" selection below). The value of 'Queue Shaper Rate' is restricted to **100 - 1000000 kbps**. If you need coarser granularity, select the '**Mbps**' unit.

Queue Shaper Rate

Controls the rate for the queue shaper. This value is restricted to **100-1000000** when the "Unit" is "**kbps**", and it is restricted to **1-3300** when the "Unit" is "**Mbps**". The default value is **500**.

Queue Shaper Unit

Controls the unit of measure for the queue shaper rate as "**kbps**" or "**Mbps**". The default value is "**kbps**". At the dropdown select the unit of measure for the Rate selected above ("**kbps**" / "**mbps**").

Queue Shaper Excess

Controls whether the queue is allowed to use excess bandwidth.

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**". The default value is "**17**".

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 LIB-44xx port.

Port Shaper Rate

Controls the rate for the port shaper. This value is restricted to **100-1000000** when the "Unit" is "**kbps**", and it is restricted to **1-13200** when the "Unit" is "**Mbps**". The default value is **500**.

Port Shaper Unit

Controls the unit of measure for the port shaper rate as "**kbps**" or "**Mbps**". The default value is "**kbps**".

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Cancel: Click to undo any changes made locally and return to the previous page.

Example

You can click the browser’s back button to display the updated “QoS Egress Port Schedulers” page.

Port	Mode	Weight							
		Q0	Q1	Q2	Q3	Q4	Q5	Q6	Q7
1	Strict Priority	-	-	-	-	-	-	-	-
2	Strict Priority	-	-	-	-	-	-	-	-
3	3 Queues Weighted	33%	33%	33%	-	-	-	-	-
4	8 Queues Weighted	13%	13%	13%	13%	13%	13%	13%	13%
5	Strict Priority	-	-	-	-	-	-	-	-
6	Strict Priority	-	-	-	-	-	-	-	-

In the screen above, LIB-4400 Ports 1 and 2 are configured for “Weighted” mode, and ports 3 and 4 are configured for “Strict Priority” mode.

Port Shaping

The **Configuration > QoS > Port Shaping** menu path displays the QoS Egress Port Shapers page. This page provides an overview of QoS Egress Port Shapers for all LIB-44xx ports.

Port	Shapers							Port	
	Q0	Q1	Q2	Q3	Q4	Q5	Q6		Q7
1	-	-	-	-	-	-	-	-	-
2	-	-	-	-	-	-	-	-	-
3	-	-	-	-	-	-	-	-	-
4	-	-	-	-	-	-	-	-	-
5	-	-	-	-	-	-	-	-	-
6	-	-	-	-	-	-	-	-	-

The displayed settings are explained below:

Port

The logical port number for the settings contained in the same row (e.g., ports 1-4 on the LIB-4400). Click on the port number to configure the port shapers for that port.

Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Shapers

Displays columns **Q0 - Q7** and **Port**.

Qn

Shows "**disabled**" or the actual port shaper rate (e.g. "**800 Mbps**").

Port

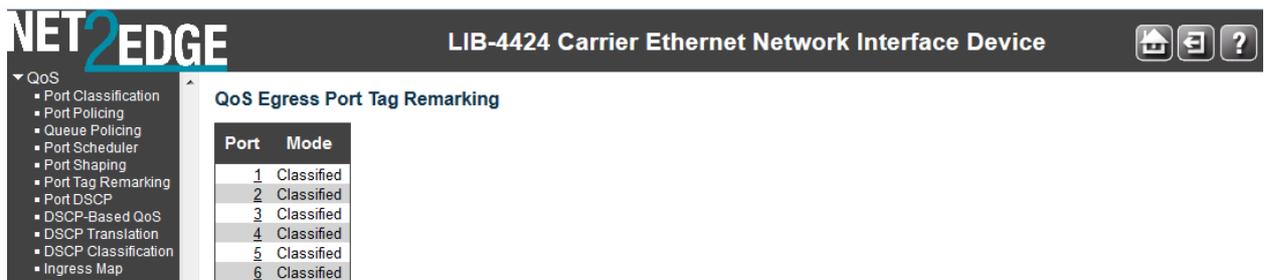
Displays the current Port rate (e.g. "**500 kbps**") or "**disabled**".

Click on a port number link in the Port column (at the far left of the table) to display that port's "QoS Egress Port Scheduler and Shapers" (described earlier in this section). The **Port** column at the far right of the table displays the current port speed.

Port Tag Remarking

The **Configuration > QoS > Port Tag Remarking** menu path displays the QoS Egress Port Remarking page.

This page provides an overview of QoS Egress Port Tag Remarking for all LIB-44xx ports.



The displayed settings are explained below:

Port

The logical port for the settings contained in the same row (e.g., ports 1-4 on the LIB-4400). Click on the port number in order to configure tag remarking. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

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 QoS class and DP level.

When you click on the port number in a row, the **QoS Egress Port Tag Remarking** page for that port displays (Port 1 in the example below).



The QoS Egress Port Tag Remarking for a specific port is configured on this page.

Tag Remarking Mode

Controls the tag remarking mode for this port.

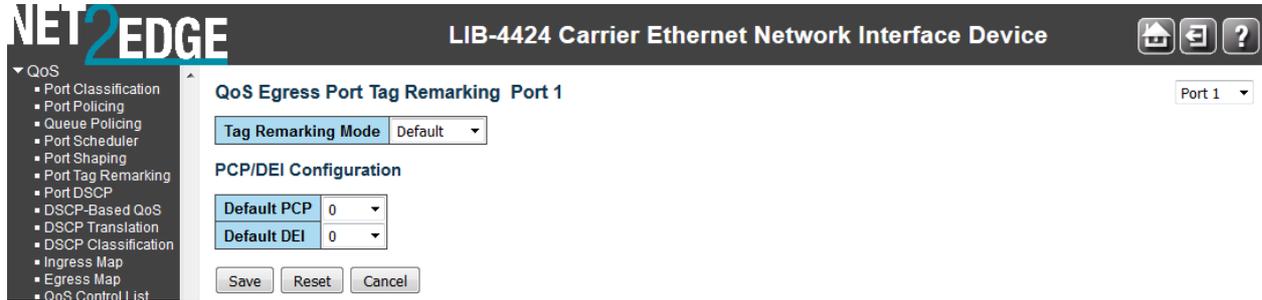
Classified: Use classified PCP/DEI values.

Default: Use default PCP/DEI values. See example below:

Mapped: Use mapped versions of QoS class and DP level. See the example below:

PCP/DEI Configuration

Controls the default PCP (0 - 7) and DEI (0, 1) values that display only when Tag Remarking Mode is set to **Default**, as shown below:



Default PCP (Priority Code Point) is a 3-bit field storing the priority level for the 802.1Q frame. It is also known as User Priority. The valid range is 0 - 7. The default is 0.

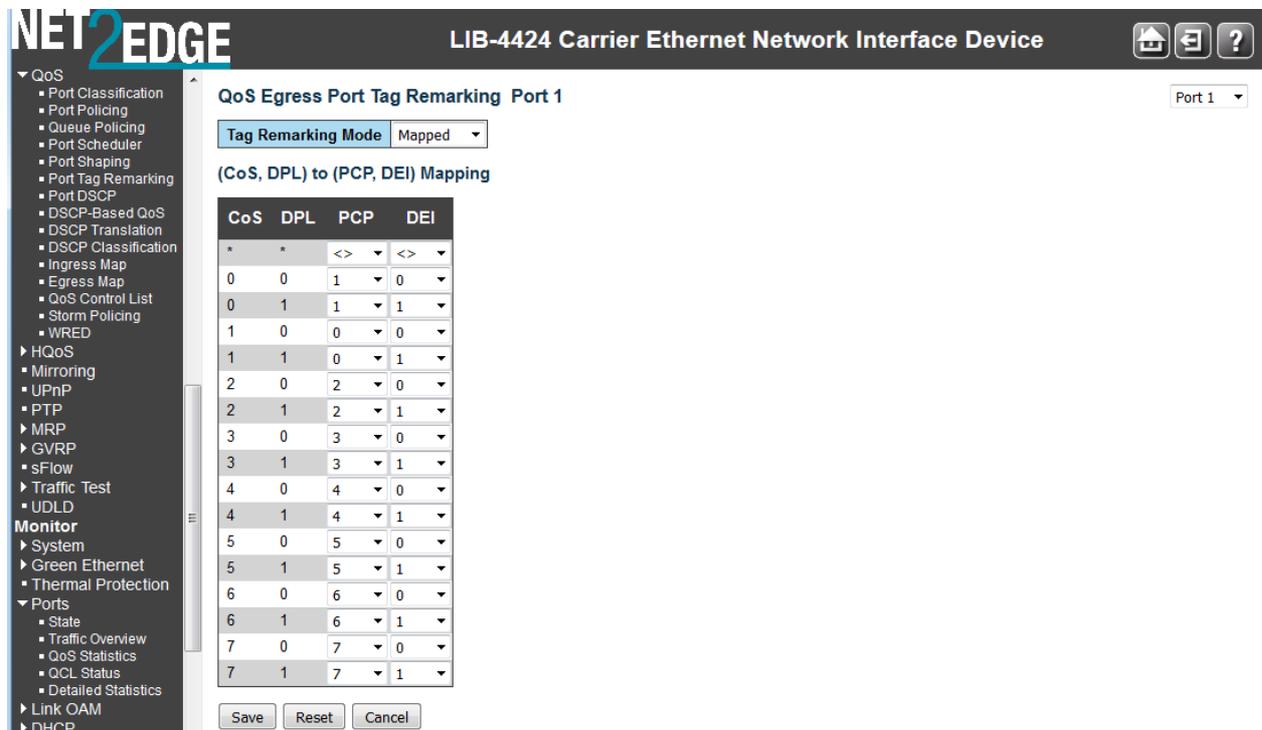
Default DEI (Drop Eligible Indicator) is a 1-bit field in the VLAN tag. The valid range is 0 - 1. The default is 0. Controls the default PCP and DEI values used when the **Tag Remarking Mode** (above) is set to **Default**.

When 0, the DEI bit in the tag is set to 0 (the default setting).

When 1, the DEI bit in the tag is set to the Classified DP level.

(QoS class, DP level) to (PCP, DEI) Mapping

Controls the mapping of the classified (QoS class, DP level) to (PCP, DEI) values when the mode is set to **Mapped** as shown below:



When **Tag Remarking Mode** is set to **Mapped**, the “(QoS class, DP level) to (PCP, DEI) Mapping” table displays.

When **Tag Remarking Mode** is set to **Default**, the “PCP/DEI Configuration” table displays. You can

click the browser's back button to display the updated "QoS Egress Port Tag Remarking" page.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Cancel: Click to undo any changes made locally and return to the previous page.

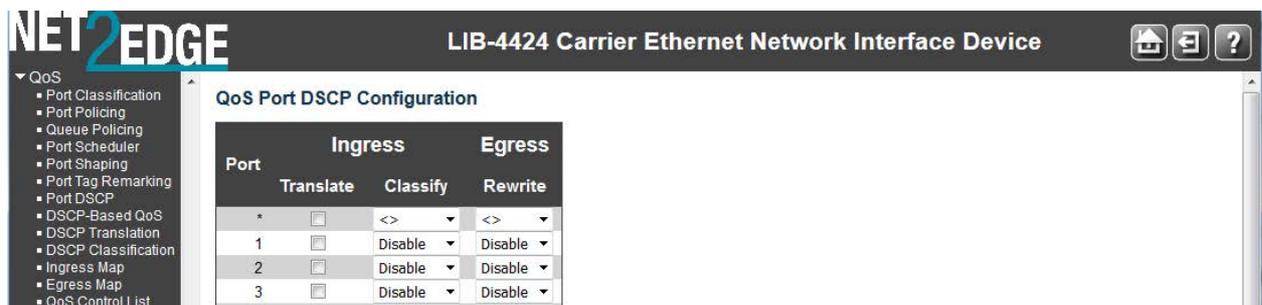
Port 1 : Use the port select box to select which port details to display.

When the **Configuration > QoS > Port Tag Remarking** configuration is done, you can view the Queuing Counters at the **Monitor > Ports > QoS Statistics** menu path.

Port DSCP

You can configure LIB-44xx QoS Port DSCP from the **Configuration > QoS > Port DSCP** menu path.

This page lets you configure the basic QoS Port DSCP Configuration settings for all LIB-44xx ports. DSCP (Differentiated Services Code Point) is a field in the header of IP packets for packet classification purposes.



The Port DSCP parameters are explained below:

Port

The Port column shows the list of ports (e.g., ports 1-4 on the LIB-4400) for which you can configure DSCP ingress and egress settings. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Ingress

In Ingress settings you can change the ingress translation and classification settings for individual ports.

There are two configuration parameters available in Ingress:

1. **Translate**
2. **Classify**

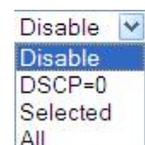
Translate

To Enable Ingress Translation click the checkbox for the port.

Classify

The Classification for a port can have one of these 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 / Rewrite

For Port Egress Rewriting select **Disable** or **Enable**, or **Remap**:

Disable: No Egress rewrite.

Enable: Rewrite enable without remapping.

Remap: DSCP from analyser is remapped and frame is remarked with remapped DSCP value.

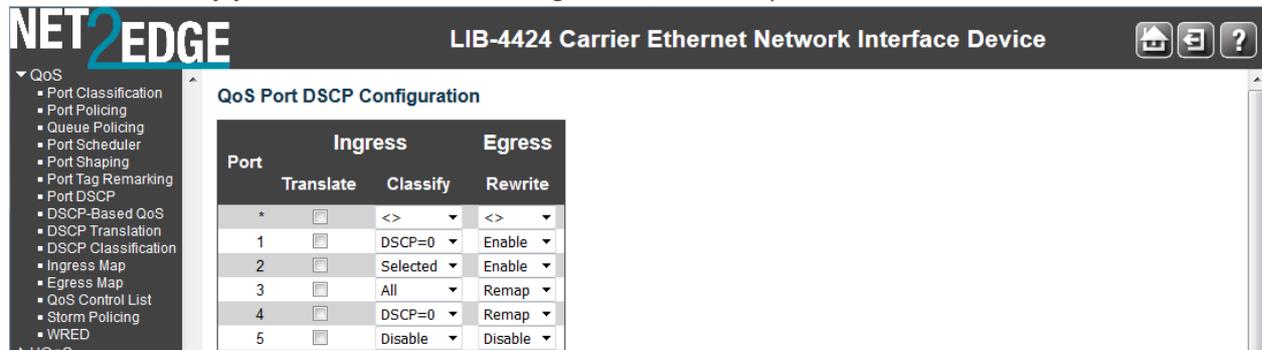


Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

When done, verify your QoS Port DSCP configuration; for example:



Click the **'Save'** button when done.

DSCP-Based QoS

The LIB-44xx DSCP-based QoS Ingress Classification page is available from the **Configuration > QoS > DSCP-Based QoS** menu path. DSCP (Differentiated Services Code Point) is a field in the header of IP packets for packet classification purposes.

3. DSCP Traffic class (IP priority): If the ingress frame is an IP packet, then the priority is decided by the IP DSCP bits. This priority is used for all frames processing inside the device. Each of the 64 traffic classes can be marked as trusted or untrusted and the QoS Class and Drop priority level can be assigned.

This page lets you configure all of the LIB-44xx QoS DSCP-based QoS Ingress Classification settings.

The screenshot shows the 'DSCP-Based QoS Ingress Classification' configuration page. On the left is a navigation menu with 'QoS' expanded to show various options. The main content area contains a table with the following data:

DSCP	Trust	CoS	DPL
*	<input type="checkbox"/>	<>	<>
0 (BE)	<input type="checkbox"/>	0	0
1	<input type="checkbox"/>	0	0
2	<input type="checkbox"/>	0	0
3	<input type="checkbox"/>	0	0
4	<input type="checkbox"/>	0	0
5	<input type="checkbox"/>	0	0
6	<input type="checkbox"/>	0	0

The displayed settings are explained below:

DSCP

The DiffServ standards define 64 DSCP values. Some of them are recommended for marking particular classes of QoS services (e.g. a value of 46 is usually recommended for real-time traffic with the strictest latency requirements).

Per the MEF, DiffServ defines several per-hop behaviours (PHBs) that provide robust QoS capabilities compared to other methods. DiffServ provides 64 different values called DiffServ Code Points, or DSCPs, that are used to determine the Class of Service (CoS). EF (Expedited Forwarding) for low delay / low loss service, AF (Assured Forwarding) in four classes for bursty real time and non-real time services, CS (Class Selector) for partial backward compatibility with IP TOS, and DF (Default Forwarding) for best effort services.

The maximum number of supported DSCP values is 64. In the DSCP column, the DSCP Name can be **BE**, **CSx**, **EFx**, or **AFx** where:

AF refers to Assured Forwarding is provided in four classes for bursty real time and non-real time services.

BE: refers to Standard (Best Effort) forwarding.

CS refers to the Class Selector (per [RFC 2474](https://www.rfc-editor.org/rfc/2474)).

EF refers to the Expedited Forwarding ([RFC 3246](https://www.rfc-editor.org/rfc/3246)). The EF PHB has the characteristics of low delay, low loss and low jitter. These characteristics are suitable for voice, video and other real-time services.

Three fundamental forwarding behaviours are defined for general use by IETF [RFC 4594](https://www.rfc-editor.org/rfc/4594): basic Default Forwarding (DF) behaviour for elastic traffic, Assured Forwarding (AF) behaviour, and Expedited Forwarding (EF) behaviour for real-time (inelastic) traffic. For additional information see the RFC at <http://tools.ietf.org/html/rfc4594>.

Trust

Check the checkbox if the DSCP value is trusted. To use DSCP requires some sort of trust between routers. Generally, DSCP values are assigned by the edge routers of an administrative domain (e.g. a Regional Network) and used by the core routers of the same domain. If there is no trust relationship between domains, the edge routers of each domain may re-assign DSCP values of ingress packets if they are already marked with a non-default (non-zero) value.

This would prevent, for example, one domain transmitting all the packets with a DSCP value indicating priority treatment into another domain and overloading that network's priority service. A

worse example could be where the DSCP value for low priority in one domain is used to mean higher priority in another. Without inspecting and re-marking the inbound DSCP values, both networks would give the opposite treatment to packets than was intended.

Some applications (e.g., VoIP or videoconferencing) mark packets by non-default DSCP values when they are generated by the application's host computer. This might be an IP phone or a videoconferencing client – although the default DSCP marking may not be consistent with the network scheme. So while this may at first seem to remove the need for an intermediate node (for example, a router) to classify the packet, the intermediate node may actually need to inspect the packets and re-mark the DSCP field.

QoS Class

Enter the QoS Class value (0-7) at the dropdown. Every incoming frame is classified to a QoS class, which is used throughout the device for providing queuing, scheduling and congestion control guarantees to the frame according to what was configured for that specific QoS class. There is a one-to-one mapping between QoS class, queue, and priority. A QoS class of 0 (zero) has the lowest priority.

DPL

Enter the Drop Precedence Level (0-3) at the dropdown. Every incoming frame is classified to a Drop Precedence Level (DP level), which is used throughout the device for providing congestion control guarantees to the frame according to what was configured for that specific DP level.

A DP level corresponds to one of three levels of 'Discard Eligible' (Yellow) frames. Packets with a DPL of 0 are least likely to be dropped and packets with a DPL of 3 are most likely to be dropped.

If congestion occurs within a class, the packets with the higher drop precedence are discarded first. To prevent issues associated with tail drop, more sophisticated drop selection algorithms such as random early detection (RED) can be used.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

DSCP Translation

The DSCP Translation page is available from the **Configuration > QoS > DSCP Translation** menu path. DSCP (Differentiated Services Code Point) is a field in the header of IP packets for packet classification purposes.

With DSCP translation, the LIB-44xx can operate as a DS boundary which can translate different domains' DSCP values. It provides classification to a remapped DSCP value or translates to a new DSCP value on ingress and also at egress it can remap to a DSCP value along with the ability to set the drop precedence level. To perform the DSCP translation, the port associated with ingress/egress also needs to enable remarking accordingly.

This page lets you configure basic QoS DSCP Translation settings. DSCP translation can be done in Ingress or Egress.

DSCP	Ingress		Egress
	Translate	Classify	Remap
*	<>	<input type="checkbox"/>	<>
0 (BE)	0 (BE)	<input type="checkbox"/>	0 (BE)
1	1	<input type="checkbox"/>	1
2	2	<input type="checkbox"/>	2
3	3	<input type="checkbox"/>	3
4	4	<input type="checkbox"/>	4
5	5	<input type="checkbox"/>	5
6	6	<input type="checkbox"/>	6

The displayed settings are explained below:

DSCP

A maximum of 64 DSCP values are supported and the valid DSCP values are **0** to **63**.

Ingress

Ingress side DSCP can be first translated to new DSCP before using the DSCP for QoS class and DPL map. There are two configuration parameters for DSCP Translation on the Ingress side:

Translate: DSCP at the Ingress side can be translated to DSCP values of **0-63**. The DSCP value can include **BE**, **CSx**, **EFx**, or **AFx** which are described below:

Classify: Click to enable Classification at the Ingress side.

The DSCP value can include **BE**, **CSx**, **EFx**, or **AFx** where:

AF refers to Assured Forwarding is provided in four classes for bursty real time and non-real time services.

BE: refers to Standard (Best Effort) forwarding.

CS refers to the Class Selector (per [RFC 2474](#)).

EF refers to the Expedited Forwarding ([RFC 3246](#)). The EF PHB has the characteristics of low delay, low loss and low jitter. These characteristics are suitable for voice, video and other real-time services.

Three fundamental forwarding behaviours are defined for general use by IETF [RFC 4594](#): basic Default Forwarding (DF) behaviour for elastic traffic, Assured Forwarding (AF) behaviour, and Expedited Forwarding (EF) behaviour for real-time (inelastic) traffic. For additional information see the RFC at <http://tools.ietf.org/html/rfc4594>.

Egress Remap

The configurable parameter is for the Egress side is **Remap**. In the **Remap** column, select the DSCP value from select menu to which you want to remap. Valid DSCP values are **0 - 63**.

The DSCP value can include **BE**, **CSx**, **EFx**, or **AFx** where:

AF refers to Assured Forwarding is provided in four classes for bursty real time and non-real time services.

BE: refers to Standard (Best Effort) forwarding.

CS refers to the Class Selector (per [RFC 2474](#)).

EF refers to the Expedited Forwarding ([RFC 3246](#)). The EF PHB has the characteristics of low delay, low loss and low jitter. These characteristics are suitable for voice, video and other real-time services.

Three fundamental forwarding behaviours are defined for general use by IETF [RFC 4594](http://tools.ietf.org/html/rfc4594): basic Default Forwarding (DF) behaviour for elastic traffic, Assured Forwarding (AF) behaviour, and Expedited Forwarding (EF) behaviour for real-time (inelastic) traffic. For additional information see the RFC at <http://tools.ietf.org/html/rfc4594>.

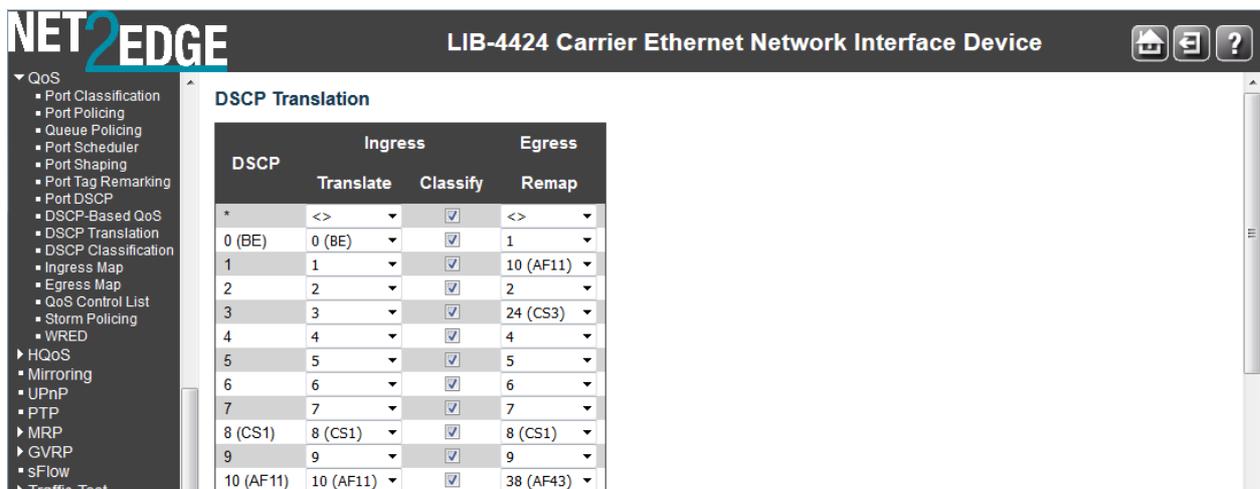
Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

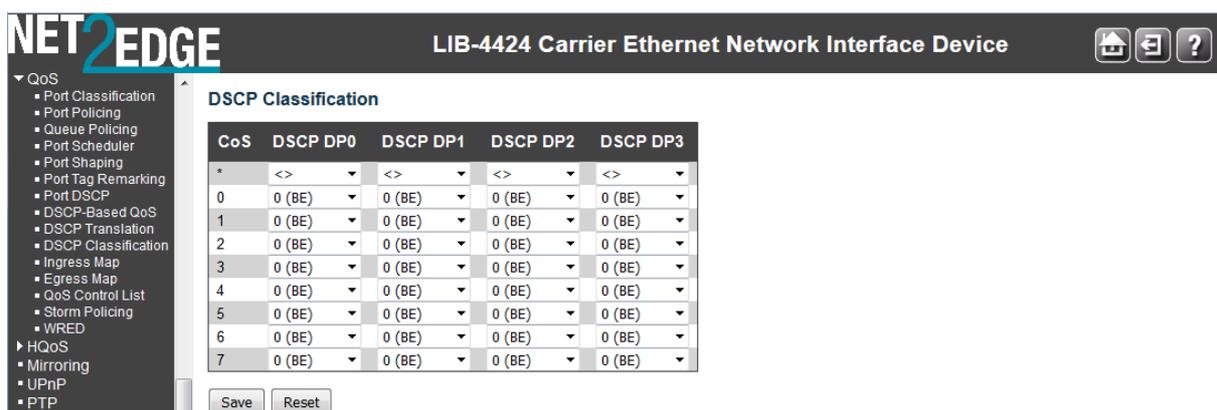
Example

Part of an edited, saved DSCP Translation page is shown below:



DSCP Classification

The DSCP Classification page is available from the **Configuration > QoS > DSCP Classification** menu path. The DSCP Classification page lets you map a DSCP value to a QoS Class and DPL value.



The displayed settings are explained below:

QoS Class

Actual QoS class values range from **0** to **7**. QoS Class (0-7) can be mapped to followed parameters. Enter the QoS Class value (**0-7**) at the dropdown. Every incoming frame is classified to a QoS class, which is used throughout the device for providing queuing, scheduling and congestion control

guarantees to the frame according to what was configured for that specific QoS class. There is a one-to-one mapping between QoS class, queue and priority. A QoS class of **0** (zero) has the lowest priority.

DSCP

Select the classified DSCP value (**0-63**) from DSCP menu to map DSCP to corresponding QoS Class and DPL values. The DiffServ standards define 64 DSCP values. Some of them are recommended for marking particular classes of QoS services (e.g., a value of 46 is usually recommended for real-time traffic with the strictest latency requirements). The maximum number of supported DSCP values is 64.

In the DSCP column, the DSCP value can include **BE**, **CSx**, **EFx**, or **AFx** where:

AF refers to Assured Forwarding.

BE: refers to Standard (Best Effort) forwarding.

CS refers to the Class Selector (per [RFC 2474](http://tools.ietf.org/html/rfc2474)).

EF refers to Expedited Forwarding ([RFC 3246](http://tools.ietf.org/html/rfc3246)).

Three fundamental forwarding behaviours are defined for general use by IETF [RFC 4594](http://tools.ietf.org/html/rfc4594): basic Default Forwarding (DF) behaviour for elastic traffic, Assured Forwarding (AF) behaviour, and Expedited Forwarding (EF) behaviour for real-time (inelastic) traffic. For additional information see the RFC at <http://tools.ietf.org/html/rfc4594>.

0 (BE)
1
2
3
4
5
6
7
8 (CS1)
9
10 (AF11)
11
12 (AF12)
13
14 (AF13)
15
16 (CS2)
17
18 (AF21)
19
20 (AF22)

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

QoS Control List

The QoS Control List page is available from the **Configuration > QoS > DSCP QoS Control List** menu path. A QCL (QoS Control List) is the list table of QCEs containing QoS control entries that classify to a specific QoS class on specific traffic objects. Each accessible traffic object contains an identifier to its QCL. The privileges determine specific traffic object to specific QoS class. A QCE (QoS Control Entry) describes the QoS class associated with a particular QCE ID.

Advanced per flow QoS: Apart from the basic QoS classification mentioned above, a QCL is provided to choose any traffic flow based on Layer 2 -4 packet headers and classify then to a particular QoS class. A complete configuration option for type of flow and the various QoS values are provided as a QCE. Some of the often used flows are VLAN based service, IP flows, TCP sessions, etc.

The QoS Control List configuration page shows the QoS Control List (QCL), which is made up of the QCEs. Each row describes a QCE that is defined. The maximum number of QCEs is 256 on each LIB-44xx.

Click on the lowest plus sign (+) to add a new QCE to the list.

The displayed settings are explained below:

QCE#

Indicates the index of QCE. The QCE (QoS Control Entry) describes the QoS class associated with a particular QCE ID.

Port

Indicates the list of ports configured with the QCE. The port members (e.g., ports 1-4 on the LIB-4400). Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Frame Type

Indicates the type of frame to look for incoming frames. Valid frame types are:

Any: The QCE will match all frame types.

Ethernet: Only Ethernet frames (with Ether Type 0x600-0xFFFF) are allowed.

LLC: Only Logical Link Control (LLC) frames are allowed.

SNAP: Only SubNetwork Access Protocol (SNAP) frames are allowed.

IPv4: The QCE will match only IPV4 frames.

IPv6: The QCE will match only IPV6 frames.

SMAC

Displays the OUI field of the Source MAC address (i.e., the first three octets (first byte) of the MAC address).

DMAC

Specify the type of Destination MAC addresses for incoming frame. Valid values are:

Any: All types of Destination MAC addresses are allowed (the default).

Unicast: Only Unicast MAC addresses are allowed.

Multicast: Only Multicast MAC addresses are allowed.

Broadcast: Only Broadcast MAC addresses are allowed.

The default value is 'Any'.

VID

Indicates the VLAN ID, either a specific VID or range of VIDs. The VID value can be **1-4094** or **Any**.

PCP

PCP (Priority Code Point): 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

DEI (Drop Eligible Indicator): Valid value of DEI can be 0, 1 or 'Any'.

Action

Indicates the classification action taken on ingress frame if parameters configured are matched with the frame's content. There are three action fields: Class, DPL and DSCP.

Class: Classified QoS Class; if a frame matches the QCE it will be put in the queue.

DPL: Drop Precedence Level; if a frame matches the QCE then the DP level will be set to the value displayed under DPL column. Every incoming frame is classified to a DP level, which is used throughout the device for providing congestion control guarantees to the frame according to what was configured for that specific DP level. A DP level of 0 corresponds to one of three levels of 'Discard Eligible' (Yellow) frames.

DSCP: If a frame matches the QCE, then DSCP will be classified with the value displayed in the DSCP column.

Modification Buttons

You can modify each QCE (QoS Control Entry) in the table using the following buttons:

- : Inserts a new QCE before the current row.
- : Edits the QCE.
- : Moves the QCE up the list.
- : Moves the QCE down the list.
- : Deletes the QCE.
- : The lowest plus sign adds a new entry at the bottom of the QCE listings.

The QCL consists of 12 QoS Control Entries (QCEs) that are searched from the top of the list to the bottom of the list for a match. The first matching QCE determines the QoS classification of the frame.

The QCE ordering is therefore important for the resulting QoS classification algorithm. If no matching QCE is found, the default QoS class is used.

The QCE (QoS Control Entry) Modification Buttons are further illustrated below:

QoS Control List Configuration

QCE#	Port	Frame Type	SMAC	DMAC	VID	PCP	DEI	Action			
								Class	DPL	DSCP	
1	1-8	Any	Any	Any	Any	Any	Any	0	Default	Default	

Insert new QCE before this QCE
 Move this QCE up
 Move this QCE down
 Delete this QCE
 Edit this QCE
 Add new QCE to end of list

Buttons

Refresh: Click to refresh the page. This will help to check the latest conflict status after releasing the resources.

When you click on the lowest plus sign (+) to add a new QCE to the end of the list, a default QCE configuration screen displays.

You can click the browser’s Back button to return to the QoS Control List page, or edit the parameters and click **Save**.

Example

Note that not all of the parameters described below will display in all situations. The **Frame Type** selection (in the **Key Parameters** table) determines which specific subset of parameters display.

The full set of QoS Control List parameters are explained below:

QCE Configuration

Port Members

Check to make this port a member, or uncheck to exclude a port, for ports 1 - 4 (on the LIB-4400). Check the checkbox to include the port in the QCL entry. By default all ports are included. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Key Parameters

Tag

Select **Any**, **Untag**, or **Tag** from the dropdown.

VID

Select **Any**, **Specific**, or **Range** from the VLAN ID (VID) dropdown. Valid value of VLAN ID can be 1-4094 or 'Any'; or, you can enter either a specific value or a range of VIDs. If you select Specific or Range, additional values are required:

Specific: enter a specific VID Value of 1-4094.

Range: enter a Range of VIDs (From: VIDa To VIDz).

PCP

Select Any, 0, 1, 2, 3, 4, 5, 6, 7, 0-1, 2-3, 4-5, 6-7, 0-3, or 4-7. This is the Priority Code Point; valid PCP values 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

Select **Any**, **0**, or **1**. This is the Drop Eligible Indicator.

SMAC

Source MAC address: 24 MS bits (OUI) or 'Any'. Select **Any** or **Specific**. If you select Specific, an entry field displays "0x" as the prefix for you to enter the specific Source MAC address.

0x

Enter a specific Source MAC in the entry field (e.g., **00-00-00**). Only displays if SMAC = Specific is selected above.

DMAC Type

Destination MAC type: possible values are Unicast (UC), Multicast (MC), Broadcast (BC) or 'Any'.

Select Any, UC, MC, or BC as the Destination MAC type, where:

Any: selects UC, MC, or BC as the Destination MAC type.

UC: selects UC (unicast) as the Destination MAC type.

MC: selects MC (multicast) as the Destination MAC type.

BC: selects BC (broadcast) as the Destination MAC type.

Frame Type

Select Any, Ethernet, LLC, SNAP, IPv4, or IPv6. Additional parameters sections display depending on the Frame Type selected here. When Frame Type =

Frame Type	Any
	Any
	Ethernet
	LLC
	SNAP
	IPv4
	IPv6

Any: Allows all types of frames. The “Action Parameters” section displays. 'Any' is the default value.

Ethernet: Ethernet Type; the valid Ethernet type can be 0x600-0xFFFF or 'Any' but excluding 0x800 (IPv4) and 0x86DD (IPv6).

LLC: selects LLC as the frame type and displays additional LLC parameters for entry (see below). This is the Logical Link Control (LLC) protocol which provides a link mechanism for upper layer protocols.

SNAP: selects the SNAP (SubNetwork Access Protocol) frame type and displays additional SNAP parameters for entry (see below).

IPv4: selects IPv4 as the frame type and displays additional IPv4 parameters for entry (see below).

IPv6: selects IPv6 as the frame type and displays additional IPv6 parameters for entry (see below).

Action Parameters

Class

For the QoS Class, select Default, 0, 1, 2, 3, 4, 5, 6, or 7.

Action Parameters

Class	0
DPL	Default
DSCP	Default

DPL

For the Drop Precedence Level (DPL) select Default, 0, or 1.

DSCP

Select Default, BE, 1, 2, 3, 4, 5, 6, 7, CS1, 9, 10, 11, 12, 13, 14, 15, CS2, 17, 18, 19, 20, 21, 22, 23, CS3, 25, 26, 27, 28, 29, 30, 31, CS4, 33, 34, 35, 36, 37, 38, 39, CS5, 41, 42, 43, 44, 45, EF, 47, CS6, 49, 50, 51, 52, 53, 54, 55, CS7, 57, 58, 59, 60, 61, 62, or 63.

Selecting 'Default' means that the default classified value is not modified by this QCE.

MAC Parameters

Ether Type

Select **Any** or **Specific**.

MAC Parameters

Ether Type	Any
------------	-----

If you select **Specific**, you must also enter a Value: 0x (e.g., Value: 0xFFFF).

MAC Parameters

Ether Type	Specific	Value: 0x	FFFF
------------	----------	-----------	------

LLC Parameters

SSAP Address

Select **Any** or **Specific**. If Specific, select a Value: 0x (e.g., Value: 0xFFFF). The SSAP Address Valid SSAP (Source Service Access Point) can vary from 0x00 to 0xFF or 'Any'.

The default value is '**Any**'.

LLC Parameters

SSAP Address	Any
DSAP Address	Any
Control	Any

DSAP Address

Select **Any** or **Specific**. If Specific, select a Value: 0x (e.g., Value: 0xFFFF). The DSAP Address Valid DSAP (Destination Service Access Point) can vary from 0x00 to 0xFF or 'Any'.

The default value is '**Any**'.

Control

Select **Any** or **Specific**. If Specific, select a Value: 0x (e.g., Value: 0xFFFF).

The valid Control field entries vary from 0x00 to 0xFF or 'Any'. The default value is '**Any**'.

SNAP Parameters

PID

Select **Any** or **Specific**. If Specific, select a Value: 0x (e.g., Value: 0xFFFF).

The Valid PID (a.k.a., Ethernet type) can have a value of 0x00-0xFFFF or 'Any'. The default is '**Any**'.

SNAP Parameters

PID	Any	▼
-----	-----	---

IPv4 Parameters

Protocol

Select **Any**, **UDP**, **TCP**, or **Other**.

If you select **UDP** or **TCP**, select Sport (Any, Specific, Range) and Dport (Any, Specific, Range). (Where 'Dport' is the Destination port and 'Sport' is the Source port).

If you select **Other**, specify a Value.

IPv4 Parameters

Protocol	Any	▼
Source IP	Any	▼
IP Fragment	Any	▼
DSCP	Any	▼

Source IP

Select **Any** or **Specific**. If you select **Specific**, select Sport and Dport (**Any**, **Specific**, **Range**). (Where 'Dport' is the Destination port and 'Sport' is the Source port). 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.

IP Fragment

Select **Any**, **Yes**, or **No**. If Range, specify the range (e.g., **BE** - **63**).

DSCP

This is the 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.

Select **Any**, **Specific**, or **Range**.

If you select **Specific**, specify the DSCP name (e.g., **BE**).

If you select **Range**, specify the DSCP range (e.g., **BE** - **63**).

IPv6 Parameters

Protocol

Select Any, **UDP**, **TCP**, or **Other**.

If you select **UDP** or **TCP**, select Sport (Any, Specific, Range) and Dport (Any, Specific, Range). (Where 'Dport' is the Destination port and 'Sport' is the Source port).

If you select **Other**, specify a Value.

IPv6 Parameters

Protocol	Any	▼
Source IP(32 LSB)	Any	▼
DSCP	Any	▼

Source IP(32 LSB)

Select **Any** or **Specific**. If you select '**Specific**', enter a 'Value' and 'Mask' (e.g., 0.0.0.0).

DSCP

Select **Any**, **Specific**, or **Range**.

If you select **Specific**, specify the DSCP name (e.g., **BE**).

If you select **Range**, specify the DSCP range (e.g., **BE** - **63**).

This is the 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.

When done editing here, click the **Save** button; the updated QoS Control List Configuration page displays with the new QCEs added to the table.

Example

QCE	Port	DMAC	SMAC	Tag Type	VID	PCP	DEI	Frame Type	CoS	DPL	Action					Ingress Map
											DSCP	PCP	DEI	Policy		
1	1-4	Unicast	00-00-00-00-00-00	Tagged	100	2-3	1	EtherType	0	Default	Default	Default	Default	Default	-	
2	5-7,18-20,25,26	Any	Any	Any	Any	Any	Any	IPv4	0	Default	Default	Default	Default	Default	-	

In the example above, two QCEs were added:

QCE #1 with “All” ports, “Ethernet” frame type, specific SMAC and DMAC, VID=100, and default Actions configured;

and

QCE #2 with ports 2-5, “IPv4” frame type, “Any” SMAC, DMAC and VID, and default Actions selected.

Use the **Modification Buttons** to Insert, Edit, Move, Delete, or Add additional QCEs as required.

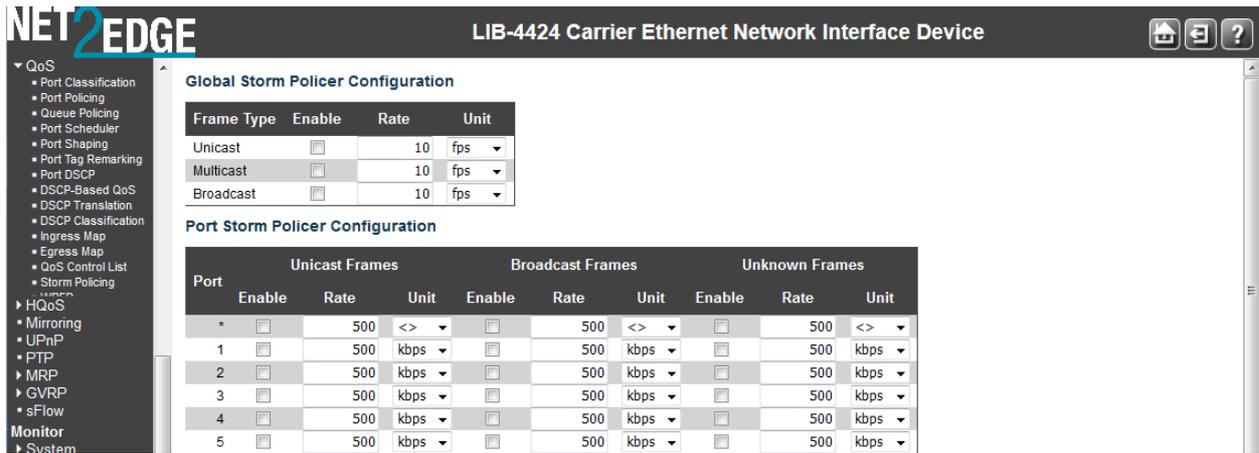
- : Inserts a new QCE before the current row.
- : Edits the QCE.
- : Moves the QCE up the list.
- : Moves the QCE down the list.
- : Deletes the QCE.
- : The lowest plus sign adds a new entry at the bottom of the QCE listings.

Storm Control

Apart from the ACL, there is device level option for storm prevention for the unicast, multicast and broadcast frames. This LIB-44xx Storm control is configured from the **Configuration > QoS > Storm Control** menu path.

This page lets you configure the storm control settings for all switch ports. There are individual storm rate controls for Unicast frames, Broadcast frames and Unknown (flooded) frames.

A traffic storm occurs when packets flood a LAN, creating excessive traffic which degrades network performance. The Storm Control feature prevents ports from being disrupted by a broadcast, multicast, or unicast traffic storm on a physical interface. Storm Control monitors incoming traffic levels over a selected traffic storm control rate and, during the interval, compares the traffic level with the selected traffic storm control level. Each port has traffic storm control levels that are used for each traffic type (Unicast, Multicast, and Broadcast). When the ingress traffic for which Storm Control is enabled reaches the traffic storm control rate you configured on the port, Storm Control drops traffic until the Storm Control interval ends. A higher rate allows more packets to pass through without dropping traffic.



The displayed settings are explained below:

Note: Frames sent to the S3280 CPU are limited to approximately 4 kbps (e.g., Management VLAN broadcasts are limited to this rate). The Management VLAN is configured at **Configuration > System > IP**.

Port

The port number for which the configuration below applies (e.g., ports 1-4 on the LIB-4400). Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Enabled

Controls whether the storm control is enabled on this switch port. The default is **disabled** (unchecked).

Rate

Controls the rate for the storm control. The default value is **500**. This value is restricted to 100-1000000 when the "Unit" is "kbps" or "fps", and it is restricted to 1-13200 when the "Unit" is "Mbps" or "kfps".

Unit

Controls the unit of measure for the storm control rate as kbps, Mbps, fps or kfps . The default value is "**kbps**".

Storm control only works on Inbound packets; it does not prevent a port from being overwhelmed with broadcasts from the core or from other access switches.

Note: Follow your organization’s policies and procedures and best practices for implementing storm control (e.g., if you should set the Multicast limit higher than the Broadcast limit, or at least set them equal, etc.).

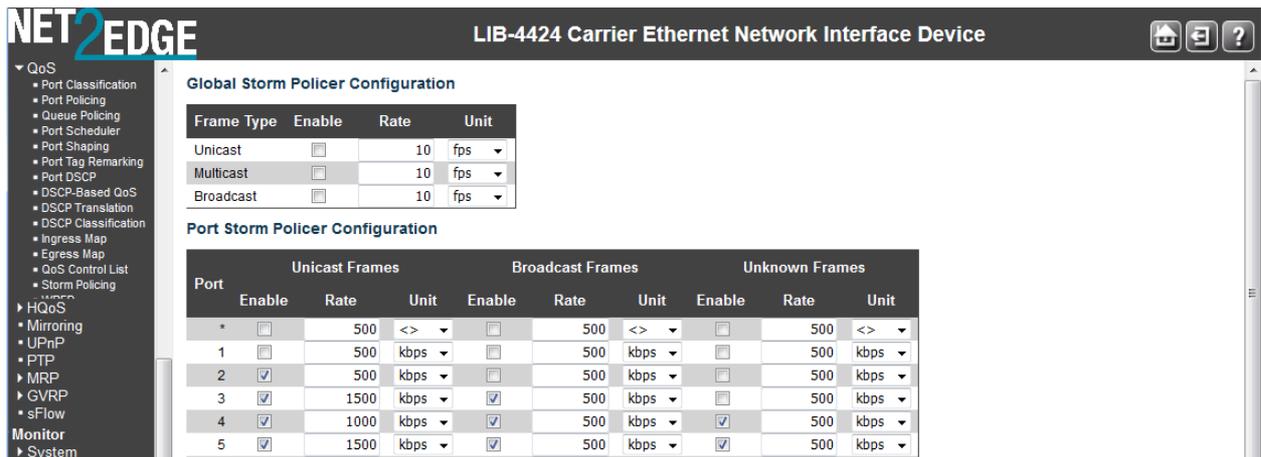
Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Example

A sample edited, saved Storm Control Configuration page is shown below:

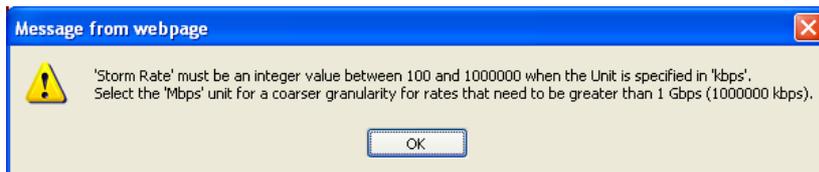


Click the 'Save' button when done configuring QoS port storm control.

Messages

Message: 'Storm Rate' must be an integer value between 100 and 1000000 when the Unit specified is 'kbps'.

Select the 'Mbps' unit for a coarser granularity for rates that need to be greater than 1 Gbps (1000000 kbps).



Meaning: You selected an unsupported combination of frame rate and unit of measure at the **Configuration > QoS > Storm Control** menu path.

Recovery: 1. Click the **OK** button to clear the webpage message. 2. Enter a valid combination of 'rate' and 'unit' parameter values.

Message: 'Storm Rate' must be an integer value between 1 and 13200 kbps.

Select the 'fps' unit for a finer granularity.

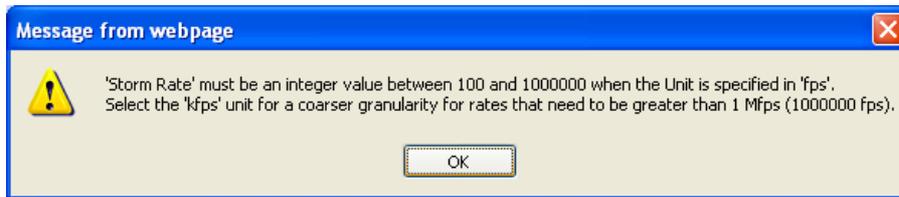


Meaning: You selected an unsupported combination of frame rate and unit of measure at the **Configuration > QoS > Storm Control** menu path.

Recovery: 1. Click the **OK** button to clear the webpage message. 2. Enter a valid combination of 'rate' and 'unit' parameter values.

Message: 'Storm Rate' must be an integer value between 100 and 1000000 when the Unit specified is 'fps'.

Select the 'kbps' unit for a coarser granularity for rates that need to be greater than 1 Mfps (1000000 kbps).



Meaning: You selected an unsupported combination of frame rate and unit of measure at the **Configuration > QoS > Storm Control** menu path.

Recovery: **1.** Click the **OK** button to clear the webpage message. **2.** Enter a valid combination of 'rate' and 'unit' parameter values.

WRED Configuration

The **Configuration > QoS > WRED** menu path lets you configure the Random Early Detection (RED) settings for queue 0 to 5. (RED cannot be applied to queue 6 and 7.) Through different RED configuration for the queues (QoS classes) it is possible to obtain Weighted Random Early Detection (WRED) operation between queues. The settings are global for all LIB-44xx ports.

Weighted Random Early Detection (WRED) is a queue management algorithm with congestion avoidance capabilities. WRED is an extension to Random Early Detection (RED) where a single queue may have several different queue thresholds, and each queue threshold is associated to a particular traffic class.

For example, a queue may have lower thresholds for lower priority packet. A queue buildup will cause the lower priority packets to be dropped, thus protecting the higher priority packets in the same queue.

This makes QoS prioritization possible for important packets from a pool of packets using the same buffer. It is more likely that standard traffic will be dropped instead of higher prioritized traffic. WRED proceeds in this order when a packet arrives:

Calculation of the average queue size.

The arriving packet is queued only if the average queue size is below the minimum queue threshold. Depending on the packet drop probability (DP), the packet is either dropped or queued if the average queue size is between the minimum and maximum queue threshold.

The packet is automatically dropped if the average queue size is greater than the maximum threshold.

The default WRED configuration page is shown below:

Group	Queue	DPL	Enable	Min	Max	Max Unit
1	0	1	<input type="checkbox"/>	0	50	Drop Probability ▼
1	0	2	<input type="checkbox"/>	0	50	Drop Probability ▼
1	0	3	<input type="checkbox"/>	0	50	Drop Probability ▼
1	1	1	<input type="checkbox"/>	0	50	Drop Probability ▼
1	1	2	<input type="checkbox"/>	0	50	Drop Probability ▼
1	1	3	<input type="checkbox"/>	0	50	Drop Probability ▼
1	2	1	<input type="checkbox"/>	0	50	Drop Probability ▼
1	2	2	<input type="checkbox"/>	0	50	Drop Probability ▼

The displayed settings are explained below:

Queue

The queue number (QoS class) for which the configuration below applies (**0-5**).

Every incoming frame is classified to a QoS class, which is used throughout the device for providing queuing, scheduling and congestion control guarantees to the frame according to what was configured for that specific QoS class. There is a one to one mapping between QoS class, queue and priority. A QoS class of **0** (zero) has the lowest priority.

Enable

Controls whether RED is enabled for this queue. Check the checkbox for a queue to enable it for WRED. The default is **disabled** (unchecked).

Min. Threshold

Controls the lower RED threshold. If the average queue filling level is below this threshold, the drop probability is zero. The valid values are **0-100**. The default is **0**.

Max. DP 1

Controls the drop probability (DP) for frames marked with Drop Precedence Level 1 (DPL 1) when the average queue filling level is 100%. The valid values are **0-100%**. The default is **1%**.

Max. DP 2

Controls the drop probability for frames marked with Drop Precedence Level 2 (DPL2) when the average queue filling level is 100%. The valid values are **0-100%**. The default is **5%**.

Max. DP 3

Controls the drop probability for frames marked with Drop Precedence Level 3 (DPL 3) when the average queue filling level is 100%. The valid values are **0-100%**. The default is **10%**.

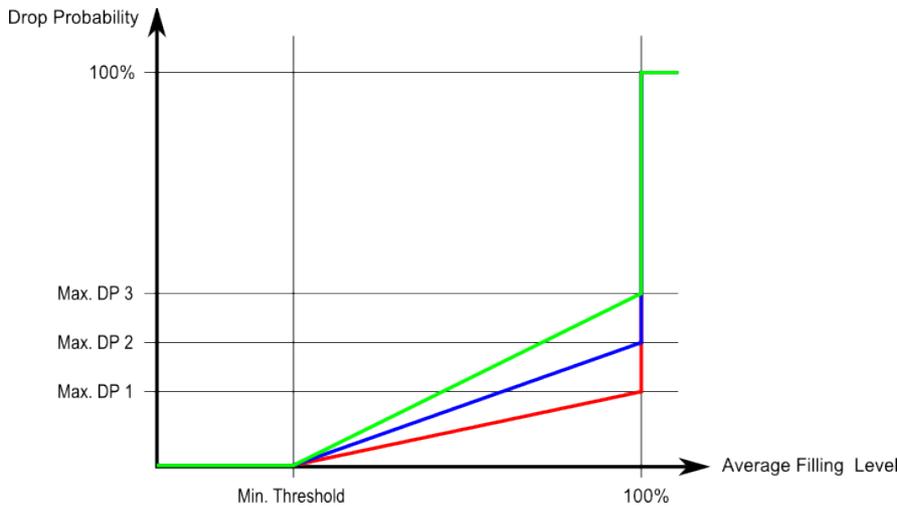
See the **Monitor > Ports > QoS Statistics** menu path for monitoring QoS Queueing Counters. RED (Random Early Detection) Drop Probability Function

The “**QoS WRED**” function lets you configure the Random Early Detection (RED) settings for queue 0 to 5. RED cannot be applied to queues 6 and 7. Through different RED configuration for the queues (QoS classes), it is possible to obtain WRED operation between queues. The settings are global for all ports in the switch.

Every incoming frame is classified to a Drop Precedence Level (DP level), which is used throughout the device for providing congestion control guarantees to the frame according to what was

configured for that specific DP level. A DP level of 0 (zero) corresponds to 'Committed' (Green) frames and a DP Level of 1 or higher corresponds to 'Discard Eligible' (Yellow) frames.

The figure below shows the drop probability function with related parameters.



Max. DP 1, **Max. DP 2** and **Max. DP 3** are the drop probabilities when the Average queue Filling Level is 100%. (Frames marked with Drop Precedence Level 0 are never dropped.)

Min. Threshold is the Average queue Filling Level where the queues randomly start dropping frames.

The drop probability for frames marked with Drop Precedence Level *n* increases linearly from zero (at Min. Threshold average queue filling level) to Max. DP *n* (at 100% Average queue Filling Level).

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

The sample screen below shows a valid, saved WRED configuration.

Group	Queue	DPL	Enable	Min	Max	Max Unit
1	0	1	<input type="checkbox"/>	0	50	Drop Probability
1	0	2	<input type="checkbox"/>	0	50	Drop Probability
1	0	3	<input checked="" type="checkbox"/>	0	50	Drop Probability
1	1	1	<input checked="" type="checkbox"/>	0	50	Drop Probability
1	1	2	<input checked="" type="checkbox"/>	0	50	Drop Probability
1	1	3	<input type="checkbox"/>	0	50	Drop Probability
1	2	1	<input type="checkbox"/>	0	50	Drop Probability
1	2	2	<input type="checkbox"/>	0	50	Drop Probability

This example shows ports 2-4 with WRED enabled at min. thresholds of 10, 20, and 30 for Ports 2-4 respectively, and using the default settings for Max. DP1, Max. DP2, and Max. DP3.

Mirroring Configuration

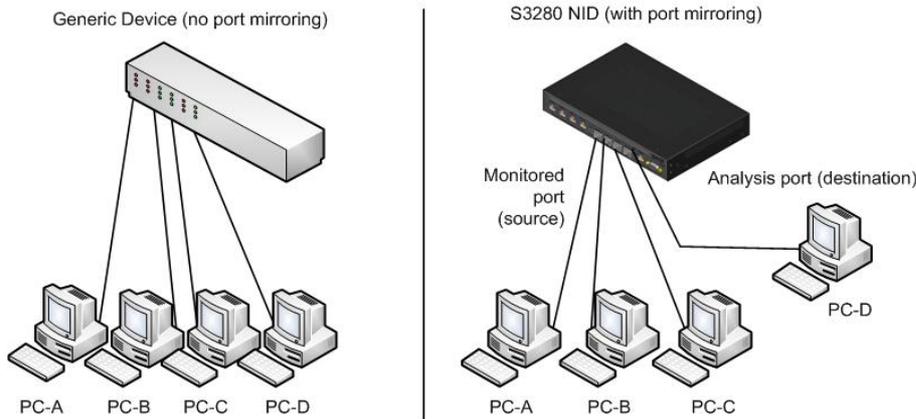
Configure port Mirroring on this page. To debug network problems, selected traffic can be copied, or mirrored, on a **mirror port** where a frame analyser can be attached to analyse the frame flow.

The traffic to be copied on the **mirror port** is selected as follows:

All frames received on a given port (also known as ingress or source mirroring).

All frames transmitted on a given port (also known as egress or destination mirroring).

The option of received or transmitted or both direction of the traffic flow can be enabled. Multiple source ports with different setting of traffic direction can be chosen to be mirrored.



Port mirror considerations:

When creating ACL rules, the traffic that falls into a particular flow can also be mirrored.

The source and mirror ports must be located on the same switch.

You can mirror the ingress or egress traffic of the source ports, or both.

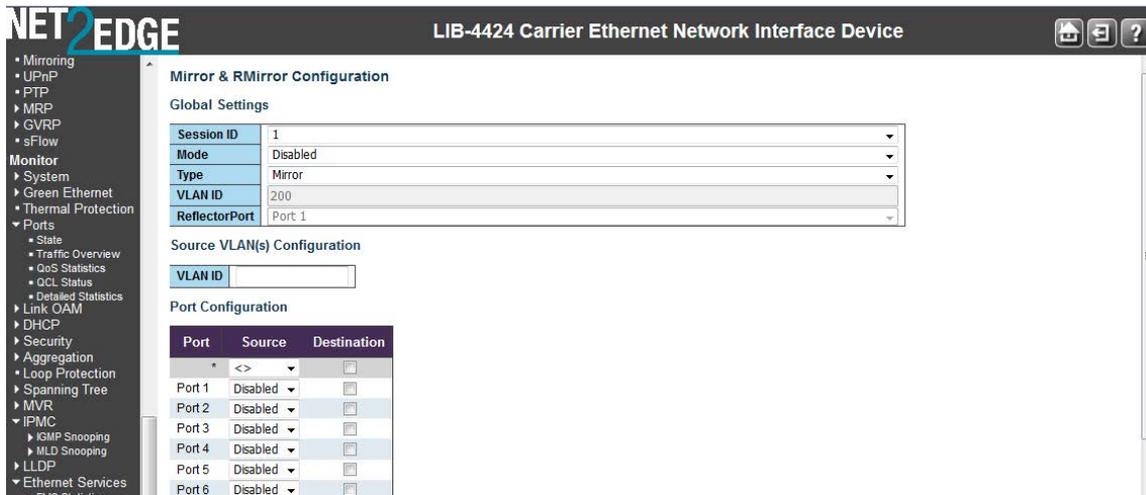
You can select more than one source port at a time. However, the more ports you mirror, the less likely the mirroring port will be able to handle all the traffic. For example, if you mirror the traffic of six heavily active ports, the destination port will likely drop packets, so it won't provide an accurate mirror of the traffic of the six source ports.

With Mirroring enabled, the mirroring port is dedicated to monitoring the traffic from the source ports, and it cannot be used for normal network operations.

Navigate to the **Configuration > Mirroring** menu path to display the Mirror Configuration page.

The default Mirror Configuration page is shown below:

Session ID	Mode	Type	VLAN ID	Reflector Port
1	Disabled	Mirror	-	-
2	Disabled	Mirror	-	-
3	Disabled	Mirror	-	-
4	Disabled	Mirror	-	-
5	Disabled	Mirror	-	-



The Mirror Configuration page parameters are explained below:

Port to mirror to

Select **Enabled** or **Disabled** for the port mirroring feature. The **Port to mirror to** is also known as the **mirror port**. Frames from ports that have either source (Rx) or destination (Tx) mirroring enabled are mirrored on this port. **Disabled** disables mirroring.

The Port / Mode table is used for Rx and Tx enabling.

Port

The logical port for the settings contained in this row (e.g., ports 1-4 on the LIB-4400). Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Mode

Select the mirror mode.

Disabled Neither frames transmitted nor frames received are mirrored (default).

Enabled Frames received and frames transmitted are mirrored on the **mirror port**.

Rx only Frames received on this port are mirrored on the **mirror port**. Frames transmitted are not mirrored.

Tx only Frames transmitted on this port are mirrored on the **mirror port**. Frames received are not mirrored.

Note: For a given port, a frame is only transmitted once. It is therefore not possible to mirror Tx frames on the **mirror port**. Because of this, the **Mode** for the selected **mirror port** is limited to **Disabled** or **Rx only**.

Buttons

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Example

An edited, saved mirror port configuration page is shown below:

NET2EDGE LIB-4424 Carrier Ethernet Network Interface Device

- Mirroring
- UPnP
- PTP
- ▶ MRP
- ▶ GVRP
- sFlow
- Monitor**
- ▶ System
- ▶ Green Ethernet
- Thermal Protection
- ▼ Ports
- State
- Traffic Overview
- QoS Statistics
- QCL Status
- Detailed Statistics
- ▶ Link OAM
- ▶ DHCP
- ▶ Security
- ▶ Aggregation
- Loop Protection
- ▶ Spanning Tree
- ▶ MVR
- ▼ IPMC
- ▶ IGMP Snooping
- ▶ MLD Snooping
- ▶ LLDP
- Ethernet Services

Mirror & RMirror Configuration

Global Settings

Session ID	1
Mode	Enabled
Type	Mirror
VLAN ID	200
ReflectorPort	Port 1

Source VLAN(s) Configuration

VLAN ID	
---------	--

Port Configuration

Port	Source	Destination
*	<>	<input type="checkbox"/>
Port 1	Disabled	<input type="checkbox"/>
Port 2	Both	<input type="checkbox"/>
Port 3	Disabled	<input type="checkbox"/>
Port 4	Disabled	<input type="checkbox"/>
Port 5	Disabled	<input checked="" type="checkbox"/>

SyncE (Synchronous Ethernet)

The **Configuration > SyncE Configuration** page lets you inspect and configure the current SyncE port settings.

Recommendation ITU-T G.8262/Y.1362 outlines requirements for timing devices used in synchronizing network equipment using synchronous Ethernet. This Recommendation defines the requirements for clocks (e.g., bandwidth, frequency accuracy, pull-in, hold-in, and pull-out ranges, noise generation, noise tolerance, noise transfer, transient response, holdover performance, etc.). Synchronous Ethernet (SyncE) as defined by ITU-T G.8261 allows for the transfer of high-quality network timing from a traceable reference, to all network elements. Because this is a physical layer process, the timing quality does not vary due to network load. Sync-E builds on the existing Ethernet standards and is backward compatible with IEEE 802.3.

The PHY devices recover the network timing from each line port and output the port recovered timing. The LIB-4400 has six clock sources (four 10G ports, one SyncE Input, and one IEEE 1588 Input), and allows each output to select recovered timing from all possible line ports.

The LIB-4424 has eight clock sources: four 10G ports, two 1G ports, one SyncE Input, and one IEEE 1588 Input and allow each output to select recovered timing from all possible line ports. If timing is compromised, the appropriate clock output can be squelched to assist with fast timing switchover.

Transmit timing is common for all ports and derived from the REFCLK, which is also used to clock the core logic. This clock is always available and is tightly controlled by the clock synchronization circuits during a timing failover. The external clock synchronization receives clocks from many possible sources, and generates a set of stable output reference clocks to be used for transmit timing. The reference clock input and output frequencies are configurable/

Synchronized Ethernet Requirements

Traditional Ethernet was originally designed to transmit asynchronous data traffic (i.e., there was no requirement to pass a synchronization signal from the source to destination). Actually, 10Base-T Ethernet is incapable of transmitting synchronization signals over the physical layer interface because a 10Base-T transmitter stops sending pulses during idle periods. A 10Base-T transmitter sends a single "I am alive" pulse every 16 ms to notify its presence to the receiving end, and these infrequent pulses do not allow clock recovery at the receiver. Idle periods in faster Ethernet (100Mbps, 1Gbps and 10Gbps) are continuously filled with pulse transitions, allowing continuous high-quality clock recovery at the receiver--good candidates for synchronized Ethernet.

Gigabit Ethernet over copper provides an additional challenge for SyncE implementation, which does not exist in Ethernet over fibre. Gigabit Ethernet over copper uses line coding as well as transmission over all four pairs of CAT-5 cable to compensate for limited bandwidth of twisted pairs used in CAT-5 cables. The transmission is done in both directions simultaneously, similar to ISDN and xDSL where digital signal processing algorithms have to be used for echo cancellation. The echo cancellation is greatly simplified if the symbol rate (frequency at which data is transmitted) is identical in both directions. This is done via the GB Ethernet Master/Slave concept. The Master generates a transmit clock locally from free-running crystal oscillator and the Slave recovers the Master clock from the received data and uses this recovered clock to transmit its own data. Master and Slave are determined during the auto-negotiation process.

Synchronization exists in Ethernet on each hop between two adjacent nodes, but it is not passed from hop to hop. Passing synchronization is fairly simple: take the recovered clock from the node receiving synchronization, and with this clock, feed all nodes that are transmitting synchronization. The recovered signal must be “cleaned” with a PLL (Phase Locked Loop) to remove jitter generated from the clock recovery circuit before being fed to the transmitting device. Also, ports must be manually set in the clock path to alternate the Master and Slave function (1000Base-T only). This is not an issue for GB Ethernet over fibre (1000Base-X or for 10 GB Ethernet (10GBASE) because one fibre is always used for transmission and the other for reception (there is no bi-directional transmission on a single fibre). So there is no need for master and slave functions.

A Gigabit or 10 Gigabit Ethernet PHY device should be able to support synchronized Ethernet as long as it provides a recovered clock on one of its output pins. The recovered clock is cleaned by the PLL and fed to the 25MHz crystal oscillator input pin on the PHY device.

It may seem that the only requirement for a PLL used in SyncE is to clean jitter from the recovered clock, which can be accomplished with general purpose PLLs. However, the PLL used in SyncE must provide additional functions beyond jitter cleaning. For example, if the receiving PHY device gets disconnected from the line, the recovered clock frequency will either stop, or it will start to “drift” depending on the clock recovery circuit implementation. A general purpose PLL will pass this large frequency change to the transmitting PHY device. As a result, not only is the transmission of synchronization signal going to fail, but the data transmission could fail as well.

The PLL used in SyncE must be able to detect failure of the recovered clock and switch the PLL to either another good reference in the system or into holdover mode. Requirements for SyncE are outlined in the timing characteristics of synchronous Ethernet equipment clock (ITU G.8262/Y1362) specifications. These specifications are based on ITU-T G.813 specification for SDH clocks. The major requirements of ITU-T G.8262/Y1362 are:

Free-run accuracy: The accuracy of PLL output when it is not driven by a reference should be equal or better than ± 4.6 ppm (part per million) over a time period of one year. This is a very accurate clock relative to the clock accuracy for traditional Ethernet (± 100 ppm).

Holdover: The PLL constantly calculates the average frequency of the locked reference. If the reference fails and no other references are available, the PLL goes into holdover mode and generates an output clock based on a calculated average value. Holdover stability depends on the resolution of the PLL averaging algorithm and the frequency stability of the oscillator used as the PLL master clock.

Reference monitoring: The PLL needs to constantly monitor the quality of its input references. If the reference deteriorates (disappears or drifts in frequency), then the PLL raises an alarm (interrupt) and switches to another valid reference.

Hitless reference switching: If the PLL's reference fails, then it will lock to another available reference without phase disturbances at its output.

Jitter and wander filtering: The PLL can be viewed as a jitter and wander filter. The narrower the loop bandwidth, the better the jitter and wander attenuation.

Jitter and wander tolerance: The PLL should tolerate large jitter and wander at its input and still maintain synchronization without raising any alarms.

These stringent requirements can only be met with a digital PLL (DPLL) like the DPLLs used for SONET/SDH clocks. The main difference is that a SyncE DPLL must be able to lock and generate clock frequencies used in Ethernet (25MHz, 125MHz, and 156.25MH) as opposed to the telecom clock frequencies used in SONET/SDH (19.44MHz, and 155.52MHz).

Configuration Considerations

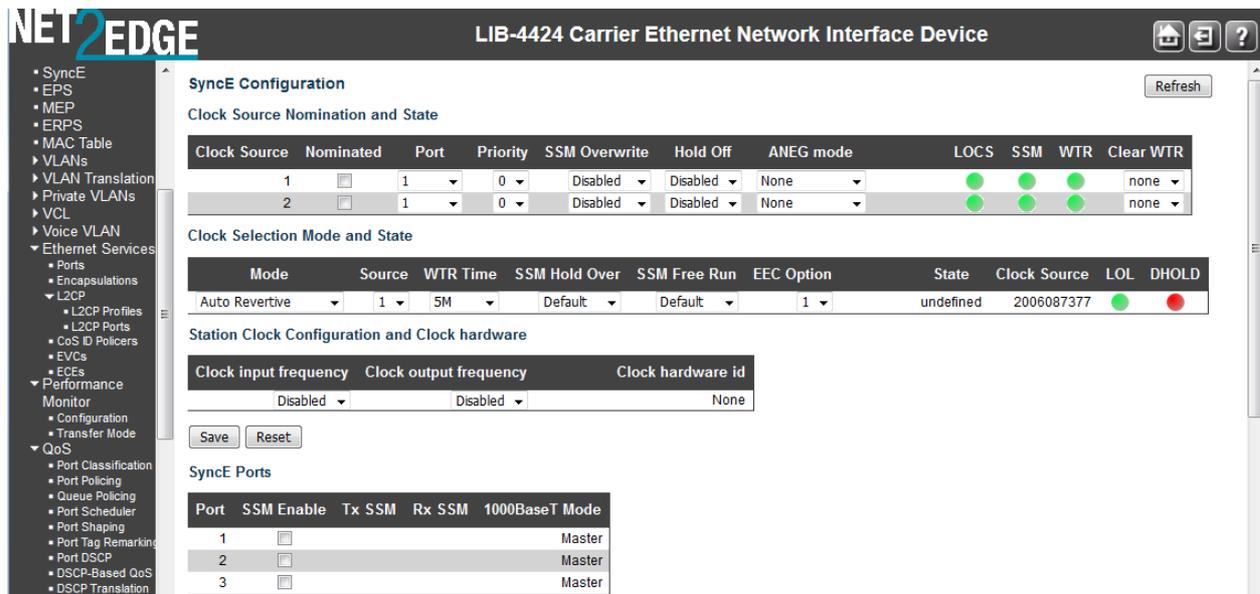
For the latest feature information and caveats, see the release notes for your particular device and software release. The prerequisites and restrictions for SyncE include:

Each network element along the synchronization path must support SyncE.

Maximum of six ports configured as clock source at a time.

Configure multiple input sources with the same priority without impacting TSM switching delay.

The default **Configuration > SyncE Configuration** page is shown below:



The SyncE parameters are described in the following sections.

Clock Source Nomination and State

These parameters can be configured for each possible clock source.

Clock Source

The instance number of the clock source. This instance number must be referenced when selecting 'Manual' Mode at the “Clock Selection Mode” field (see below).

- 1 (10G-1):** Selects clock source instance 1 as the clock source.
- 2 (10G-2):** Selects clock source instance 2 as the clock source.
- 3 (10G-3):** Selects clock source instance 3 as the clock source.
- 4 (10G-4):** Selects clock source instance 4 as the clock source.
- 5 (SyncE Input):** Selects clock source instance 5 as the clock source.
- 6 (IEEE 1588 Input):** Selects clock source instance 6 as the clock source.

Clock Source
1 (10G-1)
2 (10G-2)
3 (10G-3)
4 (10G-4)
5 (SyncE Input)
6 (IEEE 1588 Input)

Nominated

When a clock source is “nominated”, the clock output from the related PHY (Port) is enabled against the clock controller. This makes it available as a possible source in the clock selection process. The IEEE 1588 PTP Input will not be used by SyncE if it is not nominated. Check or uncheck:

- Clock Source 1 needs to be Port 1
- Clock Source 2 needs to be Port 2
- Clock Source 3 needs to be Port 3
- Clock Source 4 needs to be Port 4

Nominated
<input type="checkbox"/>

Note that in SyncE, IEEE 1588 input is disabled if no clock source is nominated.

Port

This drop down box presents the ports that are possible to select for this clock source (1-4). Indicates the port for configuration on a particular line of the Clock Source Nomination and State table.

Port	
1	▼
2	▼
3	▼
4	▼

Each port (e.g., 1-4 on the LIB-4400) represents a 10 GB SFP+ on the LIB-44xx. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Priority

The priority for this clock source (priorities 0-5). The lowest number (priority 0) is the highest priority. If two clock sources have the same priority, the lowest clock source number gets the highest priority in the clock selection process.

Priority	
0	▼
0	▼
0	▼
0	▼
0	▼
0	▼

SSM Overwrite

A selectable clock source Quality Level (QL) to overwrite any QL received in a SSM. If QL is not Received in a SSM (SSM is not enabled on this port), the SSM Overwrite QL is used as if received. The SSM Overwrite can be set to QL_NONE, indicating that the clock source is without any know quality (Lowest compared to clock source with known quality).

SSM Overwrite	
QL NONE	▼

Each SSM (Synchronization Status Message) contains a QL indication. The valid SSM Overwrite selections are:

- QL NONE:** The clock source has no known quality / No overwrite.
- QL SSUA:** The clock source is Synchronization Supply Unit A (SSUA).
- QL SSUB:** The clock source is Synchronization Supply Unit B (SSUB).
- QL EEC2:** The clock source is EEC-Option 2 (EEC synchronous Ethernet Equipment Clock) per ITU-T Rec. G.8262/Y.1362 (07/2010).
- QL EEC1:** The clock source is EEC-Option 1 (EEC synchronous Ethernet Equipment Clock) per ITU-T Rec. G.8262/Y.1362 (07/2010).

QL NONE
QL PRC
QL SSUA
QL SSUB
QL EEC2
QL EEC1

“SSM” in SyncE is an abbreviation for the Synchronization Status Message, which contains a QL indication. The “QL” is the Quality Level of a given clock source. This is received on a port in a SSM indicating the quality of the clock received in the port.

SSU (Synchronization Supply Units) are used to ensure reliable synchronisation distribution. SSU functions to:

- Filter the synchronisation signal they receive to remove the higher frequency phase noise,
- Provide distribution by providing a scalable number of outputs to synchronise other local equipment,
- Provide the ability to produce a high quality output even when their input reference is lost (Holdover Mode).

EEC Options: Recommendation ITU-T G.8262/Y.1362 contains two options for synchronous Ethernet. The first option (EEC-Option 1) applies to synchronous Ethernet equipment designed to interwork with networks optimized for the 2048-kbit/s hierarchy. These networks allow the worst-case synchronization reference chain as specified in Figure 8-5 of ITU-T G.803.

The second option (EEC-Option 2) applies to synchronous Ethernet equipment designed to interwork

with networks optimized for the 1544-kbit/s hierarchy. The synchronization reference chain for these networks is defined in clause II.3 of ITU-T G.813.

See the ITU spec for differences in terms of Frequency accuracy, Pull-in, hold-in, and pull-out ranges, Noise generation, Noise tolerance, Noise transfer, etc.

Hold Off

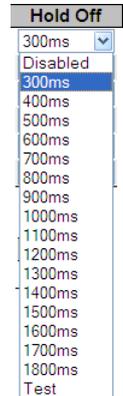
The Hold Off timer value. Active loss of clock Source will be delayed the selected amount of time. The clock selector will not change the clock source if the loss of clock condition is cleared within this time.

Disabled: Hold Off timer not enabled; Active loss of clock Source will not be delayed.

300ms – 1800ms: select a hold off timer values in 300 ms (millisecond) increments.

The valid range is 0.3 seconds (300 ms) – 1.8 seconds (1800 ms).

Test: Sets the holdoff timer to 100 in 100 ms units.



ANEG Mode

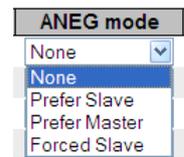
This “Auto-Negotiated” mode is only relevant for 1000BaseT ports. In order to recover clock from port, it must be negotiated to 'Slave' mode. In order to distribute clock, the port must be negotiated to 'Master' mode. These ANEG modes can be activated on a Clock Source port:

Prefer Slave: The Port will be negotiated to 'Slave' mode if possible.

Prefer Master: The Port will be negotiated to 'Master' mode if possible.

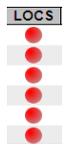
Forced Slave: The Port will be forced to 'Master' mode.

The selected port in 'Locked' state will always be negotiated to 'Slave' if possible.



LOCS

Signal is lost on this clock source (●) or signal is available on this clock source (●). Signal is lost on this clock source (Loss of Signal) indicated in red (●= Down) indicating a Loss of Clock Source (LOCS).



SSM

If SSM is enabled and not received properly (●). The type of SSM failure displays in the 'Rx SSM' field:

- = Up (green),
- = Down (red).

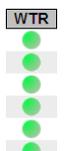


WTR

Indicate the Wait To Restore (WTR) timer state; active or inactive.

If ● displayed in green, then the Wait To Restore (WTR) timer is active (● = Up).

If ● displayed in red, then the Wait To Restore (WTR) timer (●= Down).



Clear WTR

Clears the WTR timer and makes this clock source available to the clock selection process:
clear : Clears the WTR timer and makes this clock source available to the clock selection process.

none: Makes this clock source unavailable to the clock selection process (default).



Clock Selection Mode and State

The Clock Selector is only in one instance - the one that selects between the nominated clock sources.

Mode

The definition of the 'best' clock source is firstly the one with the highest "QL" and secondly (the ones with equal QL) the highest priority. Set the Clock Selector to one of these modes:

Manual: Clock selector will select the clock source stated in Source (see below).

If this manually selected clock source is failing, the clock selector will go into holdover state.

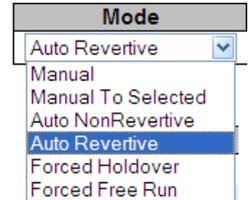
Manual To Selected: Same as Manual mode where the selected clock source will become Source.

Auto NonRevertive: Clock Selection of the best clock source is only done when the selected clock fails.

Auto Revertive: Clock Selection of the best clock source is constantly done.

Forced Hold Over: Clock Selector is forced to Free Run State.

Forced Free Run: Clock Selector is forced to Free Run State.



Source

Only relevant if "Manual" mode is selected (see above). At the dropdown, select 1-6.

The Clock selector will select this clock source (see "Mode" description above).

1 (10G-1): Selects clock source instance 1 as the clock source.

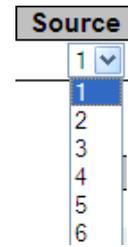
2 (10G-2): Selects clock source instance 2 as the clock source.

3 (10G-3): Selects clock source instance 3 as the clock source.

4 (10G-4): Selects clock source instance 4 as the clock source.

5 (SyncE Input): Selects clock source instance 5 as the clock source.

6 (IEEE 1588 Input): Selects clock source instance 6 as the clock source.



WTR Time

WTR is the Wait To Restore timer value in minutes. The WTR time is activated on the falling edge of a clock source failure (in Revertive mode). This means that the clock source is first available for clock selection after the WTR Time (which can be cleared). The settings are:

Disable The WTR timer is not used.

1M: The WTR time is 1 minute (60 seconds).

2M: The WTR time is 2 minutes (120 seconds).

3M: The WTR time is 3 minutes (180 seconds).

4M: The WTR time is 4 minutes (240 seconds).

5M: The WTR time is 5 minutes (300 seconds).

6M: The WTR time is 6 minutes (360 seconds).

7M: The WTR time is 7 minutes (420 seconds).

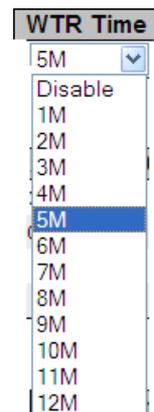
8M: The WTR time is 8 minutes (480 seconds).

9M: The WTR time is 9 minutes (540 seconds).

10M: The WTR time is 10 minutes (600 seconds).

11M: The WTR time is 11 minutes (660 seconds).

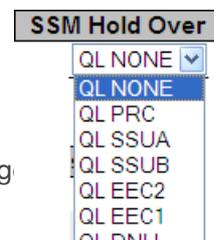
12M: The WTR time is 12 minutes (720 seconds).



SSM Hold Over

This is the transmitted SSM QL (Synchronization Status Message Quality Level) value when the clock selector is in Hold Over State. **At the dropdown, select either:**

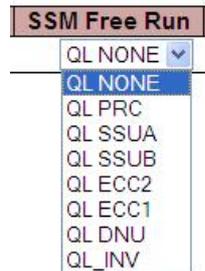
QL NONE: The clock source has no known quality / No overwrite.



- QL PRC:** The clock source is the Primary Reference Clock (PRC).
- QL SSUA:** The clock source is Synchronization Supply Unit A (SSUA).
- QL SSUB:** The clock source is Synchronization Supply Unit B (SSUB).
- QL ECC2:** The clock source is EEC-Option 2 (EEC synchronous Ethernet Equipment Clock) per ITU-T Rec. G.8262/Y.1362 (07/2010).
- QL ECC1:** The clock source is EEC-Option 1 (EEC synchronous Ethernet Equipment Clock) per ITU-T Rec. G.8262/Y.1362 (07/2010).
- QL DNU:** Do Not Use.
- QL INV:** Receiving invalid SSM (not defined) - NOT possible to set.

SSM Free Run

This is the transmitted SSM QL value when the clock selector is in Hold Over State. At the dropdown, select either QL None, QL SSUA, QLSSUB, QL ECC2, QLECC1, QL DNU, or QL_INV, where:

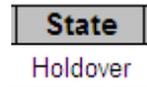


- QL NONE:** No Quality Level (QL) used.
- QL PRC:** The Primary Reference Clock (PRC) Quality Level (QL) is used.
- QL SSUA:** The SSUA (Synchronization Supply Unit A) Quality Level (QL) is used.
- QL SSUB:** The SSUB (Synchronization Supply Unit B) Quality Level (QL) is used.
- QL ECC2:** The ECC2 (EEC-Option 2) Quality Level (QL) is used.
- QL ECC1:** The ECC1 (EEC-Option 1) Quality Level (QL) is used.
- QL DNU:** The DNU (Do Not Use.) Quality Level (QL) is used.
- QL_INV:** The INV (Invalid) Quality Level (QL) is used. - NOT possible to set.

State

This indicates the current state of the clock selector. Possible selector states include:

- Free Run:** There is no external clock source to lock to (unlocked state). The Clock Selector has never been locked to a clock source long enough to calculate the hold over frequency offset to local oscillator. The frequency of this node is the frequency of the local oscillator.
- Hold Over:** There is no external clock source to lock to (unlocked state). The Clock Selector has calculated the holdover frequency offset to local oscillator. The frequency of this node is “held” to the frequency of the clock source previously locked to.
- Pre-Locked:** The Clock selector is in the process of locking to the selected clock source.
- Locked:** Clock selector is locked to the clock source indicated (see next).
- Prelocked2:** PRELOCKED2.
- Loss Of Lock:** LOSSOFLOCK2.



Clock Source

The clock source that is locked to when the clock selector is in “locked” state.



LOL

Clock selector has raised the Loss Of Lock alarm. Clock selector has raised the Loss Of Lock (LOL) alarm, where a green ● = Up, and a red ● = Down.



DHOLD

Clock selector has not yet calculated the holdover frequency offset to local oscillator. This becomes active for about 10 seconds when a new clock source is selected. Displays ● when Up (green). Displays ● when Down (red).



SyncE Ports

SyncE Ports

This section lets you configure Sync-E parameters for each possible LIB-44xx port.

Port	SSM Enable	Tx SSM	Rx SSM	1000BaseT Mode
1	<input checked="" type="checkbox"/>	QL NONE	QL LINK	Master
2	<input type="checkbox"/>			Master
3	<input type="checkbox"/>			Master
4	<input type="checkbox"/>			Master

Port

The LIB-44xx port number to configure (e.g., 1-4 on the LIB-4400). Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

SSM Enable

Enable and disable of SSM functionality on this port. The SSM (Synchronization Status Message) contains a QL (Quality Level) indication.

Tx SSM

Monitoring of the transmitted SSM QL on this port. Transmitted QL should be the Quality Level of the clock generated by this node. This means the QL of the clock source that this node is locked to.

Values can include:

QL NONE: The clock source has no known quality / No overwrite.

QL PRC: The clock source is the Primary Reference Clock (PRC).

QL SSUA: The clock source is Synchronization Supply Unit A (SSUA).

QL SSUB: The clock source is Synchronization Supply Unit B (SSUB).

QL EEC2: The clock source is EEC-Option 2 (EEC synchronous Ethernet Equipment Clock) per ITU-T Rec. G.8262/Y.1362 (07/2010).

QL EEC1: The clock source is EEC-Option 1 (EEC synchronous Ethernet Equipment Clock) per ITU-T Rec. G.8262/Y.1362 (07/2010).

QL DNU: Do Not Use.

QL INV: Receiving invalid SSM (not defined) - NOT possible to set.

Rx SSM

Monitoring of the received SSM QL (Quality Level) on this port.

If link is down on port, *QL_LINK* displays.

If no SSM (Synchronization Status Message) is received, *QL_FAIL* displays.

1000BaseT Mode

If PHY is in 1000BaseT Mode, then this is monitoring the Master/Slave mode.

Slave mode allows receive clock on a port.

Master mode allows transmit clock on a port.

External I/O Configuration

External I/O Configuration

Port	State	Frequency	Actual Frequency
SyncE Input	Enabled	8 KHz	0 Hz
SyncE Output	Enabled	64 KHz	64000 Hz

This section lets you configure LIB-44xx Sync-E Source/Direction, State, and Frequency.

Port

Indicates the SMB Port and direction.

SyncE Input: External SMB port/direction.

SyncE Output: External SMB port/direction.

State

Enable or Disable the SMB port.

Frequency

Set the Clock Frequency to one of the following values:

8 KHz: Select for 8 KHz as the frequency of the SyncE external input or output.

64 KHz: Select for 64 KHz as the frequency of the SyncE external input or output.

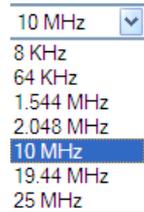
1.544 MHz: Select for 1.544 MHz as the frequency of the SyncE external input or output.

2.048 MHz: Select for 2.048 MHz as the frequency of the SyncE external input or output.

10 MHz: Select for 10 MHz as the frequency of the SyncE external input or output.

19.44 MHz: Select for 19.44 MHz as the frequency of the SyncE external input or output.

25 MHz: Select for 25 MHz as the frequency of the SyncE external input or output.



Actual Frequency

Displays the real time frequency detected into the SMB input. Active even if the SMB input is disabled (not being used internally by the device).

External I/O Options

Here you can configure the Impedance settings.

Impedance

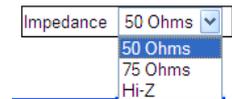
This selection box lets you select the impedance termination of the input clock. This is the value used to help match the output impedance to the transmission line impedance. Valid values are:

50 Ohms: selects 50 ohm impedance as the input clock termination.

75 Ohms: selects 75 ohm impedance as the input clock termination.

Hi-Z: not impedance termination driven, "tri-stated" or "floating".

External I/O Options



Buttons

Refresh: Click to refresh the page immediately.

Auto-refresh: Check this box to enable an automatic refresh of the page at three second intervals.

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Messages

Message: *Clock Source 1 needs to be Port 1*

Clock Source 2 needs to be Port 2

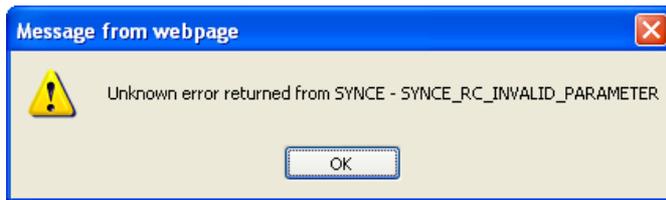


Meaning: You entered an invalid "Clock Selection Mode and State" parameter.

Recovery: 1. Click the **OK** button to clear the web page message. 2. Re-enter the parameter.

3. See "[Clock Selection Mode and State](#)" on page 328.

Message: *Unknown error returned from SYNC_E_RC_INVALID_PARAMETER*



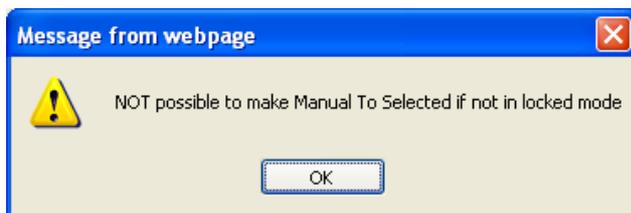
Meaning: You entered an invalid “Clock Selection Mode and State” parameter.

Recovery: 1. Click the **OK** button to clear the web page message. 2. Re-enter the parameter.

3. See “[Clock Selection Mode and State](#)” on page 328.

Message: *NOT possible to make Manual To Selected if not in Locked mode*

NOT possible to make Manuel To Selected if not in locked mode



Meaning: You entered an invalid “Clock Selection Mode and State” parameter.

Recovery: 1. Click the **OK** button to clear the web page message. 2. Re-enter the parameter.

3. See “[Clock Selection Mode and State](#)” on page 328.

Message:

Invalid parameter error returned from SYNCE

Port nominated to a clock source is already nominated

Not legal to change port on a nominated clock source - first de-nominate clock source

The selected port is not valid

Meaning: You entered an invalid “Clock Selection Mode and State” parameter.

Recovery: 1. Click the **OK** button to clear the web page message. 2. Re-enter the parameter.

3. See “[Clock Selection Mode and State](#)” on page 328.

Message:

Clock source identification 1-n (only if manual selection mode)

Clock source identification (only if manual selection mode)

manual: Selector is manually set to the chosen clock source

selected: Selector is manually set to the pt. selected clock source (not possible in unlocked mode)

selected: Selector is manually set to the selected clock source (not possible in unlocked mode)

nonrevertive: Selector is automatically selecting the best clock source - non revertively

revertive: Selector is automatically selecting the best clock source - revertively

Meaning: An issue exists with “Clock source selection mode”.

Recovery: 1. Click the **OK** button to clear the web page message. 2. Re-enter the parameter.

3. See “[Clock Selection Mode and State](#)” on page 328.

PTP Clock Configuration

You can configure LIB-44xx PTP clocking from the **Configuration > PTP** menu path. The Precision Time Protocol (PTP) is a network protocol for synchronizing computer systems' clocks.

Precise time information is especially important for distributed systems in automation technology. With PTP as described in IEEE 1588, it is possible to synchronize distributed clocks to an accuracy of less than 1 microsecond on Ethernet networks. The demands on the local clocks and the network and computing capacity are relatively low.

Two effects are evident when setting or synchronizing clocks: 1) independent clocks initially run at an "offset". To synchronize them, the less accurate clock is set to the more accurate one (offset correction).

2) real clocks do not run at exactly the same speed. Therefore, the speed of the less accurate clock has to be regulated constantly (drift correction).

PTP knows various types of clocks, and acts as a master-to-slave protocol. A clock in an end device is known as an "Ordinary" clock, and a clock in a transmission component like an Ethernet switch is a "Boundary" clock (BC) or "Transparent" clock (TC). A "Master" synchronizes the respective 'slaves' connected to it.

The synchronization process is divided into two phases. First the time difference between the master and the slave is corrected; this is the offset correction. With IEEE1588-2008, two modes are known for the synchronization process: two-step-mode and one-step-mode. The second phase of the synchronization, delay measurement, determines the run time between slave and master. It is determined by the "Delay Request" and "Delay Response" messages in a similar way, and the clocks adjusted accordingly.

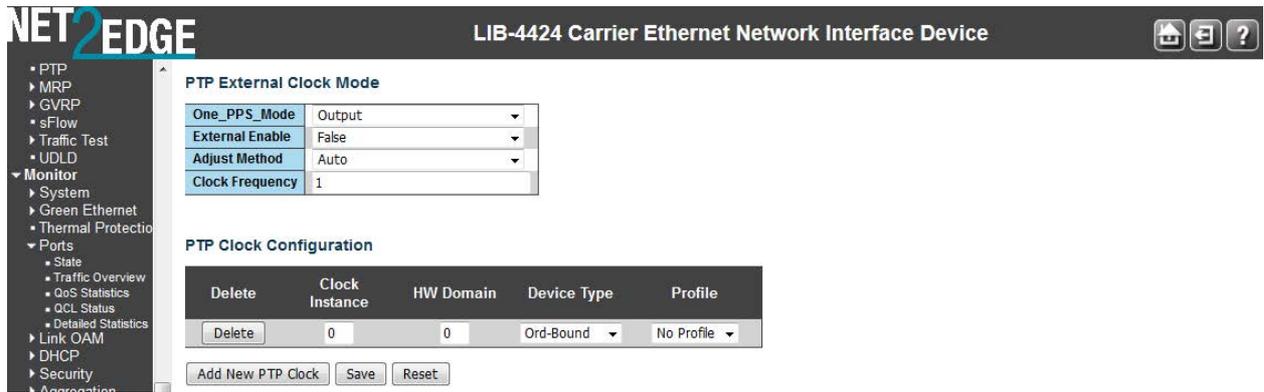
This can also be done in one-step or in two-step mode. Boundary clocks are required wherever there is a change of the communication technology or other network elements block the propagation of the PTP messages. The [IEEE 1588-2008](http://ieeexplore.ieee.org/xpl/standards.jsp) standard knows two types of transparent clocks: End-to-End (E2E) and Peer-to-Peer (P2P). See the IEEE Standards web site at <http://ieeexplore.ieee.org/xpl/standards.jsp> for current editions and amendments.

Note: at LIB-44xx v 1.0.2, PTP is available over Ethernet, IPv4 Unicast, or IPv4 Multicast only.

Note: you must have a PTP clock instance configured for accurate RFC 2544 Latency test step timestamps. PTP must be running on both devices to synchronize the Time of Day.

Note: IPv4 unicast protocol only works in Master only and Slave only clocks.

The **Configuration > PTP** menu path lets you view current PTP clock settings and configure new settings. The default PTP configuration page is shown below:



External I/O Configuration

Port

Displays the external I/O port configuration in terms of the SMB Port and direction:

IEEE 1588 Input: Lets you select Enabled or Disabled at the “State” dropdown.

IEEE 1588 Output: Lets you select Enabled or Disabled at the “State” dropdown.

State

The selection box lets you configure the External Clock output (i.e., **Enable** or **Disable** the SMB port).

The valid values are:

Enabled : Enables the port.

Disabled : Disable the port.

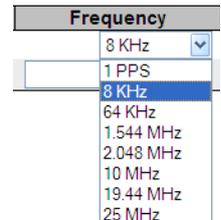
Frequency

The dropdown lets you select the Input Clock Frequency. The valid Input values are **1PPS**, **8 KHz**, **64 KHz**, **1.544 MHz**, **2.048 MHz**, **10 MHz**, **19.44 MHz**, and **25 MHz**.

1PPS, for example, is an electrical signal that has a width of less than one second and a sharply rising edge that accurately repeats once per second.

The Output entry field lets you enter a frequency in the range of 1-25,000,000 Hz.

All frequencies divisible by 4ns will be exact. All others are rounded down to the nearest frequency divisible by 4ns. The frequency box displays the frequency used internally if the input and actual do not match.



External I/O Options

Impedance

This Selection box allows you to select the impedance termination of the input clock. The selections are:

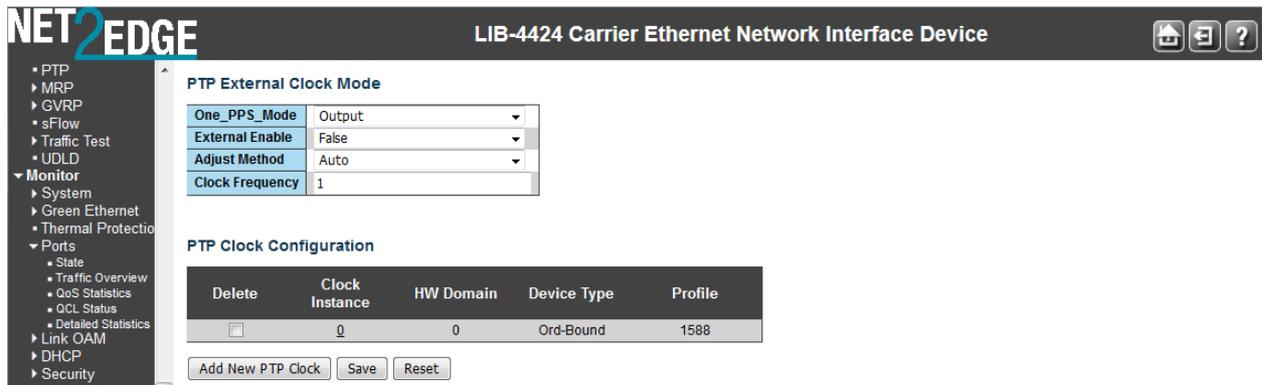
50 : Enable 50 ohm impedance.

75 : Enable 75 ohm impedance.

hiz : No impedance termination driven, "tri-stated" or "floating".

PTP Clock Configuration

By default, no clock instances are present. Click the **Add New PTP Clock** button to display the PTP configuration page shown below:



The PTP clock configuration parameters are explained below:

Delete

Click this button to immediately delete the instance. When configured and saved, check this checkbox and click 'Save' to delete the clock instance.

Clock Instance

Indicates the Instance of a particular Clock (0-3). Click on the Clock Instance number to edit its Clock details.

Device Type

Indicates the Type of the Clock Instance. Select one of the Device Types:

Inactive - no clock type is currently used. Device Type = Inactive is the state of an unused instance which may be configured but is not displayed in the "PTP Clock Configuration" table after a Save.

Ord-Bound - clock's Device Type is Ordinary-Boundary Clock.

Select Ord-Bound mode to identify the switch port that is connected to a device with the most precise clock. This is the default clock mode. The device is synchronized with the grand-master clock and operates as a parent master clock. This mode is used for switch ports when overload or heavy load conditions produce significant delay jitter.

P2p Transp - clock's Device Type is Peer to Peer Transparent Clock.

E2e Transp - clock's Device Type is End to End Transparent Clock.

Select E2e Transp mode for the switch to synchronize all switch ports with the grand master clock. The switch corrects for the delay incurred by every packet passing through it (this delay is called 'residence time'). E2e Transp mode causes less jitter and error accumulation than boundary mode.

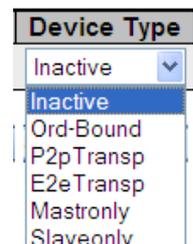
MastrOnly - clock's Device Type is Master Only.

SlaveOnly - clock's Device Type is Slave Only.

Boundary clocks and transparent clocks are mechanisms that provide accurate distribution of the PTP protocol across multi-port network components.

A boundary clock may be slaved to a master on one port and act as master on all other ports.

A transparent clock does not act as a master or slave, but instead forwards PTP event messages and provides corrections for the residence time across the bridge.



In both cases, each transparent or boundary clock includes a single PTP clock to accurately synchronize devices across the network.

Version 2 of the IEEE 1588 specification introduces transparent clocks as an alternative to implementing boundary clocks for multiport devices. The end-to-end transparent clock forwards PTP event messages, but modifies the messages for the residence time for the message to propagate from an ingress port to an egress port.

Version 2 of the IEEE 1588 specification also defines peer-to-peer transparent clocks, which measures the local link delays using the peer delay mechanism, rather than using the delay request mechanism to measure full path delay.

2 Step Flag

Static member: defined by the system, **True** if *two-step Sync* events and *Pdelay_Resp* events are used. There are two options for handling sync messages: two-step or one-step operation.

For two-step operation, both the ingress and egress timestamps need to be recorded and saved to calculate the residence time. The one-step operation may be used to eliminate timestamp transfers across the management interface.

Clock Identity

Displays the unique clock identifier (e.g., *00:c0:f2:ff:fe:56:0b:40*).

One Way

This parameter applies only to a slave. In one way mode, no delay measurements are performed (i.e., this is applicable only if frequency synchronization is needed). The master always responds to delay requests.

True: one way measurements are used.

False: one way measurements are not used.

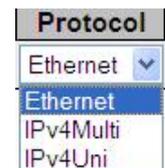
Protocol

Select the transport protocol to be used by the PTP protocol engine:

Ethernet: PTP over Ethernet multicast (the default).

IPv4Multi: PTP over IPv4 multicast.

IPv4Uni : PTP over IPv4 unicast.



Note: The IPv4 unicast protocol only works in Master-only and Slave-only clocks; see the 'Device Type' parameter description.

In a unicast Slave only clock you must also configure which master clocks to request Announce and Sync messages from. See 'Unicast Slave Configuration'.

VLAN Tag Enable

Check to enable the VLAN tagging for the PTP frames.

Note: Packets are only tagged if the port is configured for VLAN tagging (i.e., Port Type = Unaware and PortVLAN mode = None, and the port is member of the VLAN).

VID

VLAN Identifier used for tagging the PTP frames.

PCP

Priority Code Point value used for PTP frames (0-7).

Buttons

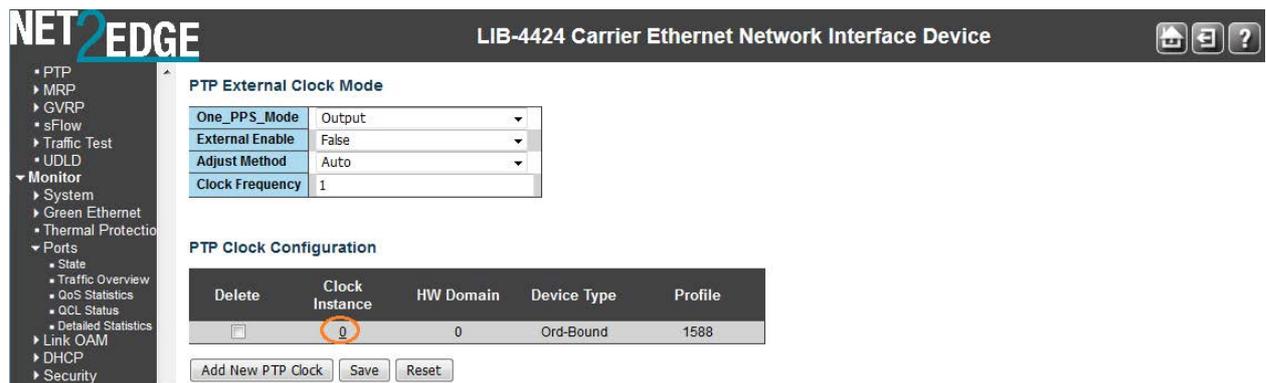
Add New PTP Clock: Click to create a new clock instance. Up to four clock instances can be created.

Save: Click to save the page immediately.

Reset: Click to reset the page immediately.

Example

The screen below shows a newly created, saved PTP Clock configuration (PTP Clock Instance = 0).



Note that no ports are configured for this PTP clock instance by default. You must define the ports list now in order to configure them in the PTP clock’s configuration (under ‘Ports Configuration’ - see below).

PTP Clock's Configuration

Click on a **Clock Instance** in the PTP Clock Configuration section at the **Configuration > PTP** menu path to display a **PTP Clock's Configuration** page with default settings.

NET2EDGE LIB-4424 Carrier Ethernet Network Interface Device

PTP Clock's Configuration and Status

Clock Type and Profile

Clock Instance	HW Domain	Device Type	Profile	Apply Profile Defaults	Filter Type
0	0	Ord-Bound	1588	Apply	ACL_BC_FULL_ON_PATH_FREQ

Port Enable and Configuration

Port Enable																													Configuration
1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	Ports Configuration
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>																											

Virtual Port Enable and Configuration

Enable	I/O Pin	Class	Accuracy	Variance	Pri1	Pri2	Local Prio
False	0	248	254	65535	128	128	128

Local Clock Current Time

PTP Time	Clock Adjustment method	Synchronize to System Clock
1970-01-08T20:29:32+00:00 000,793,940	Internal Timer	Synchronize to System Clock

Clock Current DataSet

stpRm	Offset From Master	Mean Path Delay
0	0.000,000,000	0.000,000,000

Clock Parent DataSet

Parent Port ID	Port	PStat	Var	Rate	GrandMaster ID	GrandMaster Clock Quality	Pri1	Pri2
00:01:c1:ff:fe:00:fc:b0	0	False	0	0	00:01:c1:ff:fe:00:fc:b0	Cl:248 Ac:Unknwn Va:65535	128	128

Clock Default DataSet

Device Type	One-Way	2 Step Flag	Ports	Clock Identity	Dom	Clock Quality
Ord-Bound	False	False	29	00:01:c1:ff:fe:00:fc:b0	0	Cl:248 Ac:Unknwn Va:65535

Pri1	Pri2	Local Prio	Protocol	VID	PCP	DSCP
128	128	128	Ethernet	1	0	0

Clock Time Properties DataSet

UtcOffset	Valid	leap59	leap61	Time Trac	Freq Trac	ptp Time Scale	Time Source
0	False	False	False	False	False	True	160

Leap Pending	Leap Date	Leap Type
False	1970-01-01	leap61

Save Reset

Each of the **PTP Clock's Configuration** page sections and fields are explained below:

Local Clock Current Time table

PTP Time

Shows the actual PTP time with nanosecond resolution (e.g., 2012-01-03T19:52:44+00:00 541,523,780).

Clock Adjustment method

Shows the actual clock adjustment method. The method depends on the available hardware (e.g., Software, Internal Timer, or External Timing (ET) Board).

Internal Timer: uses internal clocking only. The first clock instance configured (0) will have “Internal Timer” configured.

Software: uses software-based clocking only. Clock instances 1-3 will have “Software” configured.

hasEtBoardTiming: has an External Timing (ET) Board that is used for timing.

Synchronize to System Clock

You can click the **Synchronize to System Clock** button to synchronize the System Clock to PTP Time.

Ports Configuration

Click the link to edit the port data set for the ports assigned to this clock instance. (You must have enabled the ports list earlier in order to display the table data here now.)

When you click the [Ports Configuration](#) link, the PTP Clock's Port Data Set Configuration page displays (see description later in this section)

Clock Default DataSet Table

The clock default data set is defined in the IEEE 1588 Standard. It holds three groups of data: the static members defined at clock creation time, the Dynamic members defined by the system, and the configurable members which can be set here.

ClockId

An internal instance ID (0-3).

Type

Indicates the Type of the Clock Instance. The five Device Types are:

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.

2 Step Device Flag

Static member: defined by the system, **True** if two-step Sync events and Pdelay_Resp events are used, otherwise **False**.

Clock Identity

Displays unique clock identifier (e.g., 00:c0:f2:ff:fe:00:00:01).

Dom

The Clock domain (0-127).

Clock Quality

The clock quality is determined by the system, and includes three parts: Clock Class, Clock Accuracy and OffsetScaledLog Variance as defined in IEEE1588 (e.g., Cl:251 Ac:254 Va:65535).

The Clock Accuracy values are defined in IEEE1588 table 6 (the Clock Accuracy is set to **Unknown** by default).

ri1

Clock priority 1 (0-255) used by the BMC master select algorithm. The Best Master Clock (BMC) algorithm determines which clock is the highest quality clock within the network. The BMC (grandmaster clock) then synchronizes all other (slave) clocks in the network. If the BMC is removed from the network or is found by the BMC algorithm to no longer be the highest quality clock, the algorithm then redefines the new BMC and adjusts all other clocks accordingly. No admin input is needed.

Pri2

Clock priority 2 (0-255) used by the BMC master select algorithm.

Protocol

Transport protocol used by the PTP protocol engine:

ethernet: PTP over Ethernet multicast.

ip4multi: PTP over IPv4 multicast.

ip4uni: PTP over IPv4 unicast.

Note: IPv4 unicast protocol only works in Master only and Slave only clocks . See the parameter “Device Type”. In a unicast Slave only clock you must also configure which master clocks to request Announce and Sync messages from. See “Unicast Slave Configuration” below:

One-Way

If **True**, one way measurements are used. This parameter applies only to a slave. In one-way mode, no delay measurements are performed (i.e., this applies only if frequency synchronization is needed). The master always responds to delay requests.

VLAN Tag Enable

Enables VLAN tagging for PTP frames.

VID

VLAN Identifier used for tagging the VLAN packets.

PCP

Priority Code Point value used for PTP frames.

Clock Current DataSet table

The **Clock Current DataSet** table displays the current stpRm (e.g., **0**) Offset From Master (e.g., *0.000,000,000*) and Mean Path Delay (e.g., *0.000,000,000*) information. The clock current data set is defined in the IEEE 1588 Standard. The current data set is dynamic.

stpRm

Steps Removed: The number of PTP clocks traversed from the grandmaster to the local slave clock.

Offset from master

The time difference between the master clock and the local slave clock, measured in ns (nanoseconds).

A read only value such as 0.000,000,000).

mean Path Delay

The mean propagation time for the link between the master and the local slave. A read only value such as 0.000,000,000).

Filter Parameters

The **Filter Parameters** table displays the current DelayFilter, Period, and Distance information.

DelayFilter

The default delay filter is a low pass filter, with a time constant of $2^{**}DelayFilter*DelayRequestRate$.

Period

The default offset filter uses a minimum delay filter method (i.e., the minimum measured offset during Period samples is used in the calculation).

dist

The distance between two calculations is **Dist** periods.

If **dist** is **1** the offset is averaged over the **Period**.

If **dist** is **>1** the offset is calculated using 'min' offset.

Note: In configurations with Timestamp enabled PHYs, the period is automatically increased if $(period * dist < SyncPackets \text{ pr sec} / 4)$ (i.e., if a maximum of 4 adjustments are made pr sec).

Clock Parent DataSet table

The **Clock Parent DataSet** is defined in the IEEE 1588 standard. The parent data set is dynamic.

Parent Identity

Clock identity for the parent clock; if the local clock is not a slave, the value is the clocks own ID (e.g., `00:c0:f2:ff:fe:00:00:01`).

Port Port

The Port ID for the parent master port.

PStat

Parents Stats (always *False*).

Var

The observed parent offset scaled log variance.

Change Rate

The Observed Parent Clock Phase Change Rate (i.e., the slave clock's rate offset compared to the master). The unit is ns per sec (nanoseconds per second).

Grand Master Identity

The Clock identity for the grand master clock. If the local clock is not a slave, the value is the clock's own ID (e.g., `00:c0:f2:ff:fe:00:00:01`).

Master Clock Quality

The clock quality announced by the grand master, and includes 3 parts: Clock Class, Clock Accuracy and OffsetScaledLog Variance as defined in IEEE1588 (e.g., `Cl:251 Ac:254 Va:65535`). The Clock Accuracy values are defined in IEEE1588 - Table 6. The clock Accuracy is currently set to **Unknown** as the default.)

Pri1

Clock priority 1 announced by the grand master (e.g., `0`).

Pri2

Clock priority 2 announced by the grand master (e.g., `128`).

Clock Time Properties DataSet

The clock time properties data set is defined in the IEEE 1588 Standard. The data set is both configurable and dynamic (i.e., the parameters can be configured for a grandmaster). In a slave clock, the parameters are overwritten by the grandmasters timing properties.

Grand UtcOffset

The current UtcOffset (e.g., 0 - 10000). The UTC offset is the time offset from Coordinated Universal Time (UTC). It is typically given as hour or hour and minute. Many time zones employ two time offsets; one for standard time and one for daylight saving time.

Valid

The current Valid setting (True or False).

leap59

The current leap59 setting (True or False).

leap61

The current leap61 setting (True or False).

Time Trac

The current Time Trac setting (True or False).

Freq Trac

The current Freq Trac setting (True or False).

ptp Time Scale

The current ptp Time Scale setting (True or False).

Time Source

The valid values for the Time Source parameter are 0 - 255. Typical values include:

- 16 (0x10) ATOMIC_CLOCK
- 32 (0x20) GPS
- 48 (0x30) TERRESTRIAL_RADIO
- 64 (0x40) PTP
- 80 (0x50) NTP
- 96 (0x60) HAND_SET
- 144 (0x90) OTHER
- 160 (0xA0) INTERNAL_OSCILLATOR (the default setting)

Messages:

External Clock feature not present

External PPS feature not present

Servo Parameters

The default clock servo uses a PID regulator to calculate the current clock rate using the formula:
clockAdjustment = OffsetFromMaster/ P constant + Integral(OffsetFromMaster)/ I constant + Differential OffsetFromMaster)/ D constant

A Proportional - Integral - Derivative controller (PID controller) is a control loop feedback mechanism widely used in industrial control systems. A PID is a commonly used type of feedback controller.

A PID controller calculates an "error" value as the difference between a measured process variable and a desired setpoint. The PID controller tries to minimize the error by adjusting the process control inputs.

The PID controller calculation involves three separate constant parameters: the Proportional, the Integral, and the Derivative values (**P**, **I**, and **D**). These values can be interpreted in terms of time, where:

P depends on the present error,

I depends on the accumulation of past errors, and

D is a prediction of future errors, based on current rate of change.

The weighted sum of these three actions is used to adjust the process via a control element. The Proportional, Integral, and Derivative terms are summed to calculate the output of the PID controller. By tuning these three parameters in the PID controller algorithm, the controller can provide control action designed for specific process requirements.

The Servo Parameters are explained below:

Display

If **True** then *Offset From Master*, *MeanPathDelay* and *clockAdjustment* are logged on the debug terminal.

P-enable

If **True** the **P** part of the algorithm (**Proportional**) is included in the calculation.

I-Enable

If **True** the **I** part of the algorithm (**Integral**) is included in the calculation.

D-enable

If **True** the **D** part of the algorithm (**Derivative**) is included in the calculation.

'P' constant

The **Proportional** value [1-1000]. The **Proportional** value makes a change to the output that is proportional to the current error value.

'I' constant

The **Integral** value [1-10000]. The **Integral** is the sum of the instantaneous error over time and gives the accumulated offset that should have been corrected previously.

'D' constant

The **Derivative** value [1-10000]. The **Derivative** of the process error is calculated by determining the slope of the error over time and multiplying this rate of change by the derivative gain. The derivative term slows the rate of change of the controller output.

Unicast Slave Configuration

When operating in IPv4 Unicast mode, the slave is configured up to five master IP addresses. The slave then requests *Announce* messages from all the configured masters. The slave uses the BMC algorithm to select one as the master clock; the slave then requests *Sync* messages from the selected master.

Duration

The number of seconds a master is requested to send *Announce*/*Sync* messages. The request is repeated from the slave each $\text{Duration}/4$ seconds.

ip_address

The IPv4 Address of the Master clock.

grant

The granted repetition period for the sync message.

CommState

The state of the communication with the master, possible values are:

IDLE : The entry is not in use.

INIT : *Announce* is sent to the master (Waiting for a response).

CONN : The master has responded.

SELL : The assigned master is selected as current master.

SYNC : The master is sending *Sync* messages.

Buttons

Refresh: Click to refresh the page immediately.

Auto-refresh: Check this box to enable an automatic refresh of the page at three second intervals.

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

PTP Clock's Port Data Set Configuration

Click on the [Ports Configuration Page](#) link in the PTP Clock's Configuration section to display the **PTP Clock's Port Data Set Configuration** table. The port data set is defined in the IEEE 1588 Standard.

It has three groups of data: the static members, the dynamic members, and configurable members which can be set here.

The ports displayed here are the ports enable (checkboxes checked) at the PTP Clock Configuration table in the Port List column (e.g., Ports 2, 3 and 4 in the sample screen below).

Port	Stat	MDR	PeerMeanPathDel	Anv	ATo	Syv	Dlm	MPR	Delay Asymmetry	Ingress Latency	Egress Latency	Version	Mcast Addr	Not Slave	Local Prio	2 Step Flag
1	dsbl	3	0.000,000,000	1	3	0	e2e	0	0	0	0	2	Default	False	128	Clock Def.
2	dsbl	3	0.000,000,000	1	3	0	e2e	0	0	0	0	2	Default	False	128	Clock Def.

The related **The PTP Clock's Port Data Set Configuration** table parameters are explained below:

Port

Static member port Identity : Port number [1-max port no.]. This is the LIB-44xx port count (i.e., **1-4** on the LIB-4400).

Stat

Dynamic member *portState*: The current state of the port. The clock's port status (e.g., **dsbl** or **p2pt**).

This is the current state of the port.

MDR

Dynamic member log *Min Delay Req Interval*: The delay request interval announced by the master (e.g., **0** or **3**).

Peer Mean Path Del

The path delay measured by the port in P2P mode. In E2E mode this value is **0** (e.g., **0.000,000,000**).

Anv

The interval for issuing announce messages in master state (e.g., **1**). The valid range is **-3** to **4**.

ATo

The timeout for receiving announce messages on the port (e.g., **3**). The valid range is **-1** to **10**.

Syv

The interval for issuing sync messages in the master (e.g., **0**). The Sync Interval must be an integer value between **-7** and **4**.

Dlm

Configurable member *delayMechanism*: the delay mechanism used for the port (e.g., **p2pt** for Peer to Peer Transparent), where:

e2e: End to end delay measurement.

p2pt: Peer to peer delay measurement.

dsbl: Delay is disabled.

Dlm can be defined per port in an Ordinary/Boundary clock. In a transparent clock, all ports use the same delay mechanism, as determined by the clock type.

MPR

The interval for issuing *Delay Req* messages for the port in E2e mode. This value is announced from the master to the slave in an announce message. The value is reflected in the MDR field in the Slave.

This is the interval for issuing *Pdelay_Req* messages for the port in P2P mode (e.g., 0 or 3). The valid range is -7 to 5.

Note: This value is interpreted as an interval (i.e., MPR = 0 => 1 *Delay_Req pr sec*) independent of the Sync rate (e.g., 3).

Delay Asymmetry

If the transmission delay for a link is not symmetric, the asymmetry can be configured here (e.g., 0.000,000,000). See IEEE 1588 Section 7.4.2 Communication path asymmetry.

The valid range is -1000000 to 1000000.

Ingress latency

Ingress latency measured in ns (nanoseconds), as defined in IEEE 1588 Section 7.3.4.2.

The valid range is -1000000 to 1000000.

Egress Latency

Egress latency measured in ns (nanoseconds), as defined in IEEE 1588 Section 7.3.4.2.

The valid range is -1000000 to 1000000.

Version

The current implementation supports PTP version 2 (e.g., version 2) only.

Buttons

Auto-refresh: Check this box to refresh the page automatically every 3 seconds.

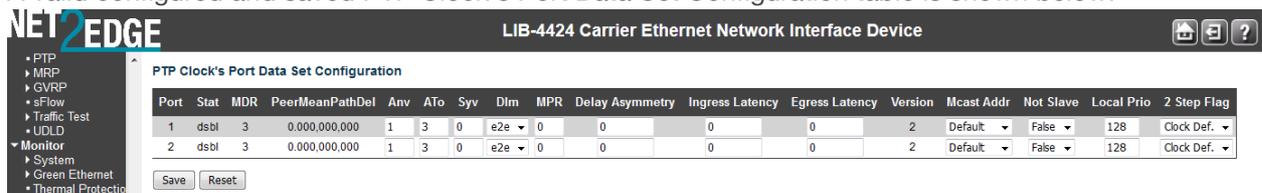
Refresh: Click to refresh the page; any changes made locally will be undone.

Save: Click to save changes.

Reset: Click to undo any changes made locally and revert to previously saved values.

Example

A valid configured and saved PTP Clock's Port Data Set Configuration table is shown below:



Service Activation

The **Configuration > Service Activation** menu path lets you define SA system, profiles, and test configuration parameters.

See the EtherSAT User Guide manual for EtherSAT / RFC 2544 configuration and operation.

Monitor Main Menu

The **Monitor** main menu lets you view and track LIB-44xx operating functions. (The related operating functions are defined at the **Configuration** main menu path.)

Each of the Monitor sub-menu functions is described below:

Monitor > System > Information

LIB-44xx system information is displayed at the **Monitor > System > Information** menu path.

(System information is entered from the

[Configuration > System](#) menu path.)

The screenshot shows the NET2EDGE web interface for a LIB-4424 Carrier Ethernet Network Interface Device. The main content area is titled "System Information" and contains the following data:

System	
Contact	
Name	LIB-4424
Location	Basingstoke
Hardware	
MAC Address	00-01-c1-00-fc-b0
Chip ID	VSC7468
Time	
System Date	2016-12-03T10:19:18+00:00
System Uptime	7d 20:39:10
Software	
Software Version	LIB-4424dev-build by jon@jon-VirtualBox 2016-11-17T14:44:45+00:00 Config:net2edge_ce_jr2_24 SDK:v02.25-smb
Software Date	2016-11-17T14:44:45+00:00
Code Revision	e666a26+
Acknowledgments	Details

The system information parameters are explained below:

Contact

Displays the system contact configured at **Configuration > System > Information > System Contact**.

Name

Displays the system name configured at **Configuration > System > Information > System Name**.

Location

Displays the system location configured at **Configuration > System > Information > System Location**.

MAC Address

Displays the MAC Address of this LIB-44xx (e.g., *00-c0-f2-56-08-b0*).

Chip ID

The Chip ID of this LIB-44xx (e.g., *VSC7428 Rev. C*).

System Date

The current (GMT) system time and date (e.g., *1970-01-01T19:05:44+00:00*). The system time is obtained through the configured timing server, if any is configured.

System Uptime

The period of time the device has been operational (e.g., *2d 19:05:44* or 2 days, 19 hours, five minutes, and 44 seconds).

Software Version

The software version of this LIB-44xx (e.g., *LIB-4400 (standalone) 1.9.4*).

Software Date

The date and time when the LIB-44xx software was produced (e.g., *2013-08-20T08:33:41-05:00*).

Acknowledgements

Click the [Details](#) link to display the related open source components. See “[Appendix B - Licenses](#)” on page [495](#).

Buttons

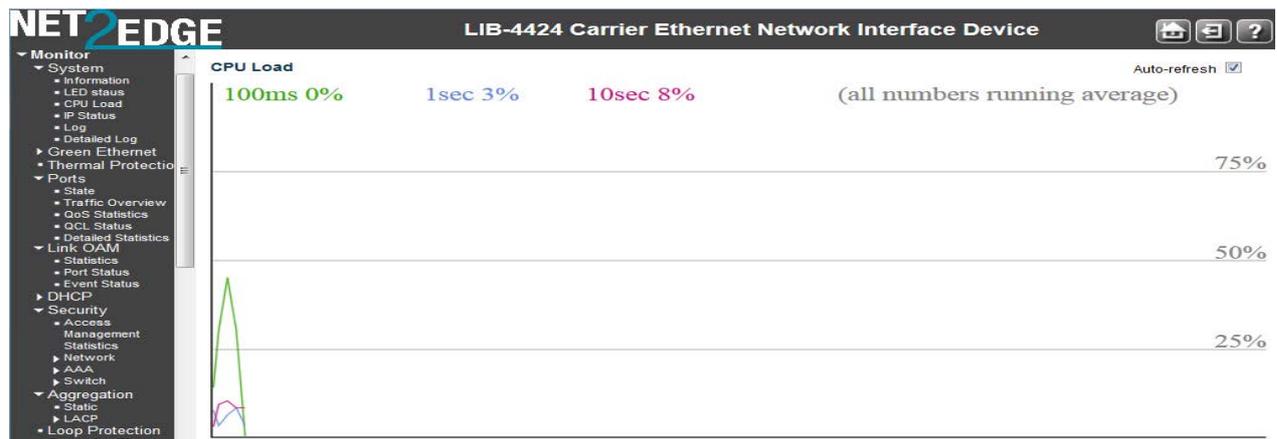
Auto-refresh: Check this checkbox to enable an automatic refresh of the page every three seconds.

Refresh: Click to refresh the page; any changes made locally will be undone.

Monitor > System > CPU Load

This page displays the CPU load, using an SVG graph. The load is measured as averaged over the last 100 milliseconds, 1 second and 10 seconds intervals. The last 120 samples are graphed, and the last numbers are displayed as text as well.

By default the message “Microsoft Internet Explorer needs the [Adobe SVG Plugin](#) to display this page.” Displays.



Your browser must support the SVG format in order to display the SVG graph. Consult the [SVG Wiki](#) for more information on browser support. Specifically, at the time of writing, Microsoft Internet Explorer must have a plugin installed to support SVG.

Download Adobe® SVG Viewer 3 for data-driven, interactive SVG graphics on the web.
Read the [Release Notes](#) and support documentation for important information about this release.
Note that Adobe [announced](#) discontinued support for Adobe SVG Viewer on January 1, 2009.

Scalable Vector Graphics (SVG) is a set of specifications for an XML-based file format for describing two-dimensional vector graphics, both static (interactive) and dynamic (animated). The SVG specification is an open standard that has been under development by the W3C since 1999. SVG images and their behaviours are defined in XML text files. This means that they can be searched, indexed, scripted and, if required, compressed. All major modern web browsers have at least some degree of SVG support, and can render SVG markup directly, including Mozilla Firefox, Internet Explorer 9, Google Chrome, Opera, and Safari. However, versions of Microsoft Internet Explorer before IE9 support SVG natively.

To download, install and run the Adobe SVG Plugin

If you want to download, install, and run the plugin, perform the steps below: (As an alternative, you can just run the plugin without downloading/installing, as explained in the next section.)

Click the [Adobe SVG Plugin](#) link. A new window opens with the Adobe SVG Viewer download area. To install the Adobe SVG Viewer, double-click the downloaded installer and follow the on-screen instructions.

Click **Save** at the prompt. Click **Run** to run the program. (You must have Admin privileges on your computer.)

If you are not using Internet Explorer, you may need to restart your browser before viewing SVG.

To run the Adobe SVG Plugin:

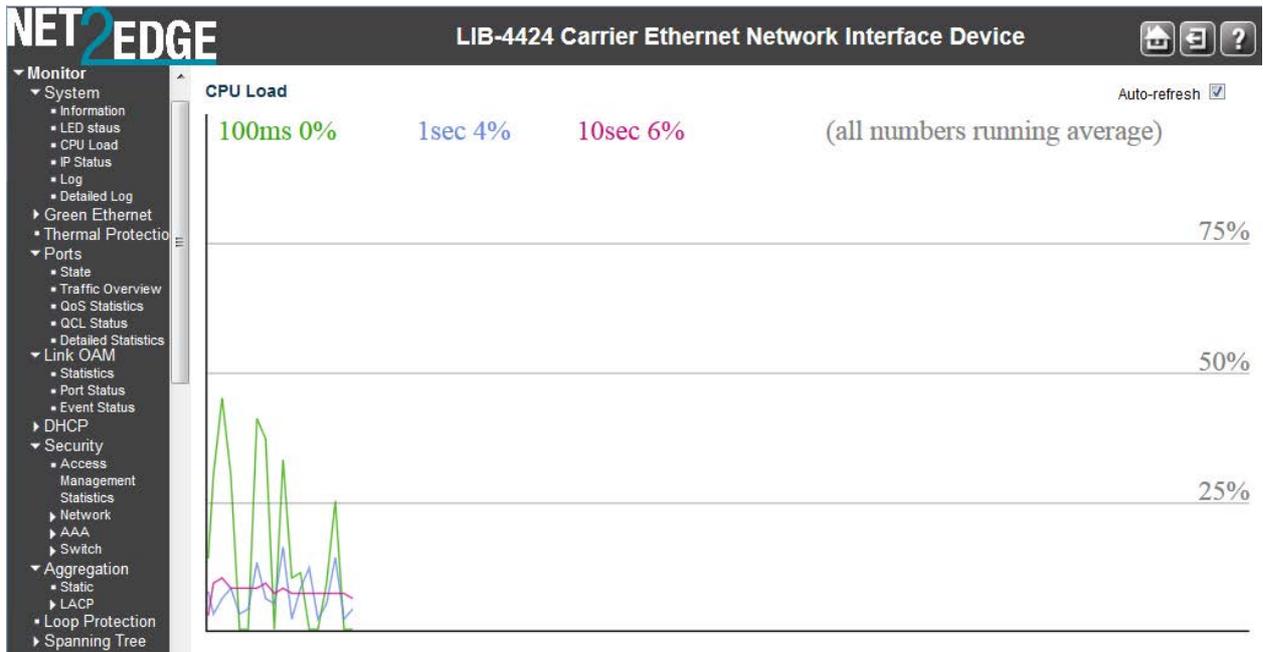
If you want to run the plugin without installing it, perform the steps below:

At the information / message bar above the LIB-44xx web page, click *“This website wants to run the following add-on: ‘SVG Viewer 3.0.2 for Netscape’ from ‘Adobe Systems, ’ (unverified publisher). If you trust the website and the add-on, and want to allow it to run, click here ...*

The CPU Load page displays using an SVG graph. The message *“Collecting data, please wait ...”* displays momentarily, and then some initial CPU load data displays.

The load is measured as averaged over the last 100 milliseconds, 1 second and 10 seconds intervals. The last 120 samples are graphed, and the last numbers are displayed as text as well.

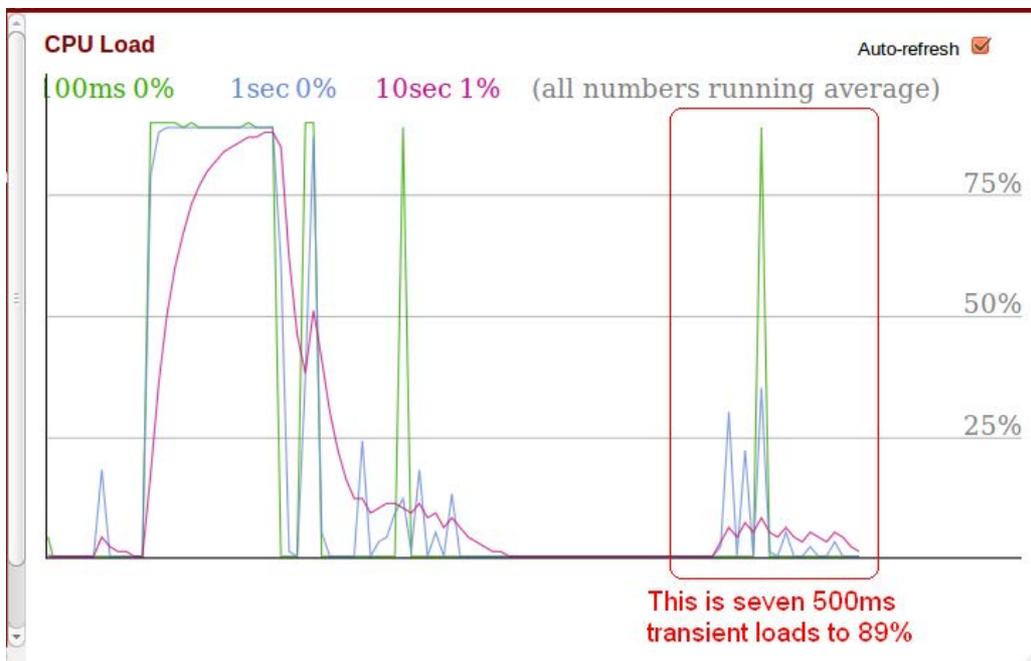
If you let it run for a while, the CPU Load graph will look something like this:



Buttons

Auto-refresh: Check this checkbox to enable an automatic refresh of the page at 3 second intervals.

The screen below shows a workload introduced at 500 ms intervals that creates transient spikes in the graphs to 89% load.



For troubleshooting High CPU utilization conditions, see “[Troubleshooting High CPU Load Conditions](#)” on page 433.

Monitor > System > Log

The **Monitor > System > Log** menu path displays the System Log Information page. The LIB-44xx system log information is provided here. (System Logging is configured from the **Configuration > System > Log** menu path.)

Syslog is a method to collect messages from devices to a server running a syslog **daemon**. Logging to a central syslog server helps in aggregation of logs and alerts which is useful for troubleshooting.

NET2EDGE LIB-4424 Carrier Ethernet Network Interface Device

System Log Information

Level: All
Clear Level: All

The total number of entries is 42 for the given level.

Start from ID 1 with 20 entries per page.

ID	Level	Time	Message
1	Warning	2016-11-25T13:40:46+00:00	illegal version[0x5] at section 0, expected: 0x6
2	Warning	2016-11-25T13:40:46+00:00	illegal version[0x5] at section 1, expected: 0x6
3	Informational	2016-11-25T13:40:47+00:00	SYS-BOOTING: Switch just made a cold boot.
4	Notice	2016-11-25T13:40:51+00:00	LINK-UPDOWN: Interface Vlan 1, changed state to down.
5	Notice	2016-11-25T13:40:51+00:00	LINK-UPDOWN: Interface Vlan 1, changed state to down.
6	Notice	2016-11-25T13:40:51+00:00	LINK-UPDOWN: Interface GigabitEthernet 1/21, changed state to up.
7	Notice	2016-11-25T13:40:57+00:00	LINK-UPDOWN: Interface Vlan 1, changed state to up.
8	Notice	2016-11-25T13:41:00+00:00	LINK-CHANGED: Interface GigabitEthernet 1/1, changed state to up (MEP).
9	Notice	2016-11-25T13:41:00+00:00	LINK-CHANGED: Interface GigabitEthernet 1/2, changed state to up (MEP).
10	Notice	2016-11-25T13:41:00+00:00	LINK-CHANGED: Interface GigabitEthernet 1/3, changed state to up (MEP).
11	Notice	2016-11-25T13:41:00+00:00	LINK-CHANGED: Interface GigabitEthernet 1/4, changed state to up (MEP).
12	Notice	2016-11-25T13:41:00+00:00	LINK-CHANGED: Interface GigabitEthernet 1/5, changed state to up (MEP).
13	Notice	2016-11-25T13:41:00+00:00	LINK-CHANGED: Interface GigabitEthernet 1/6, changed state to up (MEP).
14	Notice	2016-11-25T13:41:00+00:00	LINK-CHANGED: Interface GigabitEthernet 1/7, changed state to up (MEP).
15	Notice	2016-11-25T13:41:00+00:00	LINK-CHANGED: Interface GigabitEthernet 1/8, changed state to up (MEP).
16	Notice	2016-11-25T13:41:00+00:00	LINK-CHANGED: Interface 2.5GigabitEthernet 1/1, changed state to up (MEP).
17	Notice	2016-11-25T13:41:00+00:00	LINK-CHANGED: Interface 2.5GigabitEthernet 1/2, changed state to up (MEP).
18	Notice	2016-11-25T13:41:00+00:00	LINK-CHANGED: Interface 2.5GigabitEthernet 1/3, changed state to up (MEP).
19	Notice	2016-11-25T13:41:00+00:00	LINK-CHANGED: Interface 2.5GigabitEthernet 1/4, changed state to up (MEP).
20	Notice	2016-11-25T13:41:00+00:00	LINK-CHANGED: Interface GigabitEthernet 1/9, changed state to up (MEP).

For detailed syslog information, click an ID number in one of the lines. The LIB-44xx system log information is explained below:

ID

The ID of the system log entry. Each level can display up to 12 entries. Each ID is hot linked to its details page (see “Detailed Log” below).

Level

The level of the system log entry. The following level types are supported:

Info: Information level of the system log. Normal operational messages - used for reporting, measuring throughput, etc. - require no action.

Warning: Warning level of the system log. Warning messages - not an error, but indication that an error will occur if action is not taken (e.g. *file system 85% full*). Each item must be resolved within a given time.

Error: Error level of the system log. Non-urgent failures - these should be relayed to developers or admins. Each item must be resolved within a given time.

All: All three levels of information are logged (Info, Warning, and Error).

Clear Level

The level of the system log entries to be cleared (**Info**, **Warning**, **Error**, or **All**).

Time

The time of the system log entry. The format is *yyyy-mm-ddThh:mm:ss+<offset>*. For example: "1970-01-03T18:15:35+00:00".

Message

The message of the system log entry. If the selected level has no syslog info to report, the message "No system log entries" displays. See "System Log Messages" on page 488 for more information.

Buttons

Auto-refresh: Check this checkbox to enable an automatic refresh of the page at 3 second intervals.

Refresh: Updates the system log entries, starting from the current entry ID. **Note:** Repeated Refresh button action causes configuration display changes with each refresh action (e.g., displays "level" = warning, "clear level"= all, press F5/Refresh; "clear level"= "error"; press F5; the configuration goes back to "clear level"= all).

Clear: Flushes all system log entries.

|<<: First page; updates the system log entries, starting from the first available entry ID.

<<: Previous page; updates the system log entries, ending at the last entry currently displayed.

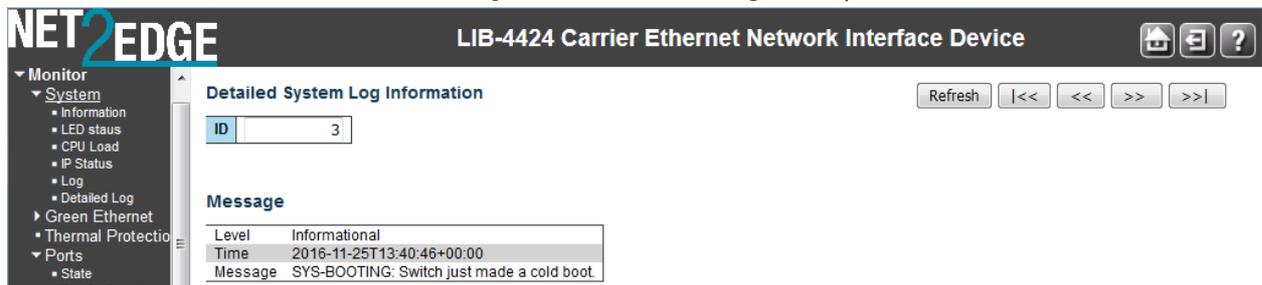
>>: Next page; updates the system log entries, starting from the last entry currently displayed.

>>|: Last page; updates the system log entries, ending at the last available entry ID.

See "System Log Messages" on page 488 for more information.

Detailed System Log Information

You can access the syslog details either by clicking on an ID column number in the System Log Information table, or via the **Monitor > System > Detailed Log** menu path.



Use the browser’s back button to return to the **Monitor > System > Log** page.

System Log Message Summary

The System Log information is summarized below in terms of syslog level, message, and description.

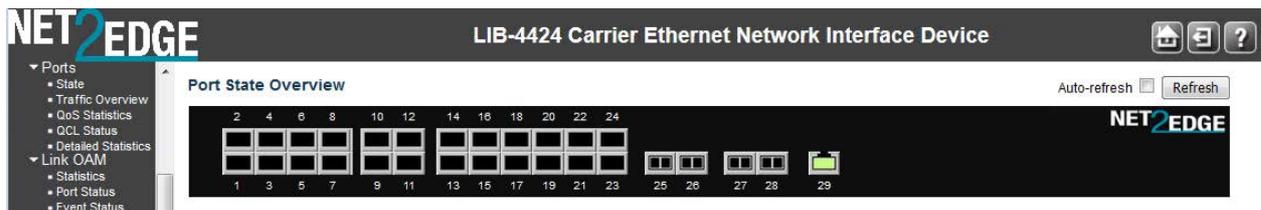
Level	Sample Message	Description
Info	Switch just made a cold boot.	The LIB-44xx was restarted. See "Maintenance > Restart Device".
	Link up on port x	The most recent link status on the port x is 'link up'. Port Link is up - no action needed.
	Link down on port x	The most recent link status on the port x is 'link down'. See Configuration > Ports > Port Configuration.
	Using primary power source.	Normal power on operation.
	Frame of 243 bytes received on port 1MAC	Normal frame size information.

	<i>Port 1 shut down</i>	Normal port shutdown information.
Warning	<i>E api/cil 17:42:26 29/126_action_check#6036:</i>	ACL policer and EVC policer cannot both be enabled. Disable one or the other.
	<i>Device running outside operating temperature: -128</i>	The PS module temp should be checked.
Error	<i>E api/cil 17:42:45 29/126_acl_policer_free#6069:</i>	Error: policer 0 already free. The EVC policer or
	<i>VLAN Port Configuration Ingress Filter Conflict - MSTP</i>	Verify the Ingress Policers, Port Policing, or Queue Policing configuration. See the Configuration > Security > Network > ACL or the Configuration > QoS menu path.
	<i>VLAN Port Configuration Ingress Filter Conflict - ERPS</i>	Verify the ERPS configuration at the Configuration > ERPS menu path.

See “[System Log Messages](#)” on page 488 for System Log message descriptions.

Monitor > Ports > State

From the **Monitor > Ports > State** menu path you can view the Port State Overview table. This page provides an overview of the current LIB-44xx port states (this page is also the default (startup) page).



The LIB-4400 port states are shown and explained below:

Port	State		
	Disabled	Down	Link
PORT 1 - PORT 4 (SFP ports)			
PORT 5 (MGMT)			

The /LIB-4424 port states are shown and explained below:

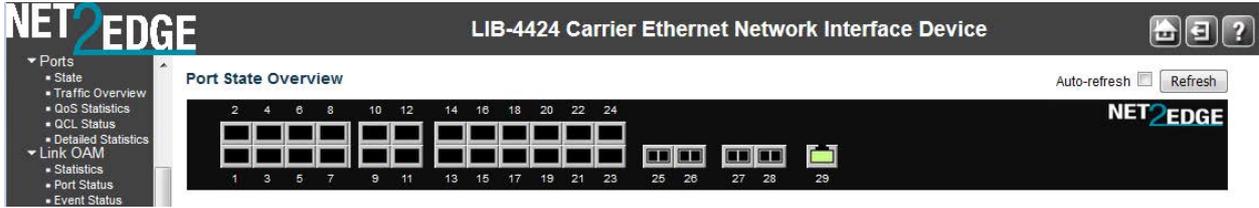
Port	State		
	Disabled	Down	Link
PORT 1 - PORT 28 (SFP ports)			
MGMT port			

Buttons

Auto-refresh: Check this box to refresh the page automatically every three seconds.

Refresh: Click to refresh the page; any changes made locally will be undone.

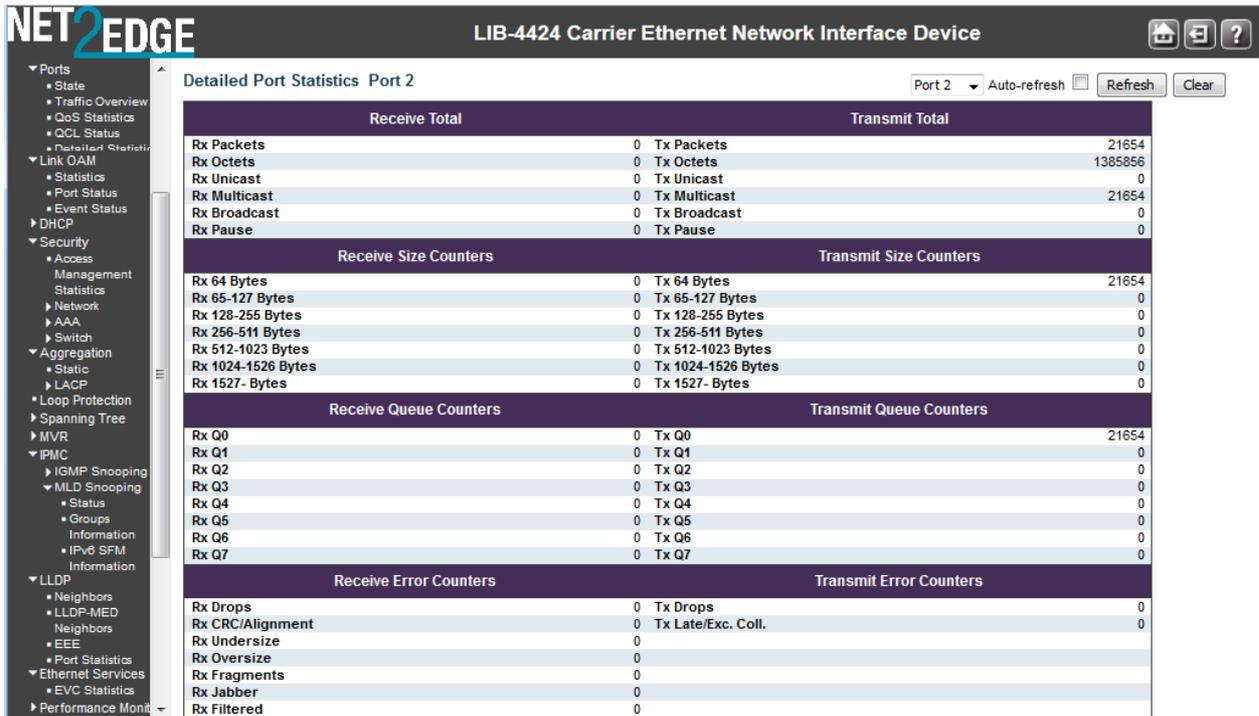
Example: the cursor over a port for the speed of that port (e.g., *Port 1: Down* or *Port 2: Down*, or *Port 29: 1Gfdx*). For example:



Detailed Port Statistics

Left mouse click on a port to display that port’s ‘Detailed Port Statistics’ page. (You can also reach this page from the **Monitor > Ports > Detailed Statistics** menu path.) See “

[Detailed Port Statistics](#)” on page 317 for more information on the ‘Detailed Port Statistics’ parameters. A sample ‘Detailed Port Statistics’ page is shown below (for Port 1).



These Port 1 ‘Detailed Port Statistics’ parameters are expanded below:

Detailed Port Statistics Port 1

Port 1 Auto-refresh

Receive Total		Transmit Total	
Rx Packets	0	Tx Packets	21665
Rx Octets	0	Tx Octets	1386560
Rx Unicast	0	Tx Unicast	0
Rx Multicast	0	Tx Multicast	21665
Rx Broadcast	0	Tx Broadcast	0
Rx Pause	0	Tx Pause	0
Receive Size Counters		Transmit Size Counters	
Rx 64 Bytes	0	Tx 64 Bytes	21665
Rx 65-127 Bytes	0	Tx 65-127 Bytes	0
Rx 128-255 Bytes	0	Tx 128-255 Bytes	0
Rx 256-511 Bytes	0	Tx 256-511 Bytes	0
Rx 512-1023 Bytes	0	Tx 512-1023 Bytes	0
Rx 1024-1526 Bytes	0	Tx 1024-1526 Bytes	0
Rx 1527- Bytes	0	Tx 1527- Bytes	0
Receive Queue Counters		Transmit Queue Counters	
Rx Q0	0	Tx Q0	21665
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	0
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	0		

Buttons

The port select box determines which port is affected by clicking the buttons.

Refresh: Click to refresh the page immediately.

Clear: Clears the counters for the selected port.

Auto-refresh: Check this box to enable an automatic refresh of the page at 3 second intervals.

Monitor > Ports > Traffic Overview

From the Monitor > Ports > Traffic Overview menu path you can view the Port Statistics Overview table. This page provides an overview of general traffic statistics for all LIB-44xx ports.

Port	Packets		Bytes		Errors		Drops		Filtered
	Received	Transmitted	Received	Transmitted	Received	Transmitted	Received	Transmitted	Received
1	0	41703	0	2668992	0	0	0	0	0
2	0	41703	0	2668992	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

The displayed counters are explained below:

Port

The logical port for the settings contained in the same row (e.g., ports 1-4 on the LIB-4400). Provides a link to the Detailed Port Statistics page for each port. Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

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.

Buttons

Refresh: Click to refresh the page immediately.

Clear: Clears the counters for all ports.

Auto-refresh: Check this checkbox to enable an automatic refresh every three (3) seconds.

Detailed Port Statistics

Left mouse click on a port to display that port's 'Detailed Port Statistics' page. (You can also reach this page from the **Monitor > Ports > Detailed Statistics** menu path.) See ‘

Detailed Port Statistics” on page 317 for more information on the ‘Detailed Port Statistics’ parameters.

Example

The screen below shows Port Statistics reporting Port 3 filtering traffic:

Port Statistics Overview

Auto-refresh Refresh Clear

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	687757	146	44016448	17931	0	0	0	0	687757
4	0	0	0	0	0	0	0	0	0
5	0	0	0	0	0	0	0	0	0

Monitor > Ports > QoS Statistics

From the **Monitor > Ports > QoS Statistics** menu path you can view the Queuing Counters page. This page provides statistics for the various queues for all LIB-44xx ports.

The LIB-44xx supports eight transmission queues per port based on user priority of each frame. This gives you the option to change the output queue mapping based on priority. Statistics for each output queue on each port are available. The option to restrict a port to a specific number of MAC entries to be learned is provided to allow the provider to restrict number of devices that can be serviced by this port.

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	41708	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
2	0	41708	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	0
4	0	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	0
6	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

The displayed QoS Statistics counters are explained below:

Port

The logical port for the settings contained in the same row. Click on an individual port number in the ‘Port’ column to display the Detailed Port Statistics for that port (e.g., ports 1-4 on the LIB-4400).

Note that the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Qn

There are eight QoS queues per port. Q0 is the lowest priority queue.

Rx

The number of received packets per queue.

Tx

The number of transmitted packets per queue.

Buttons

Auto-refresh: Click to refresh the page immediately.

Refresh: Clears the counters for all ports.

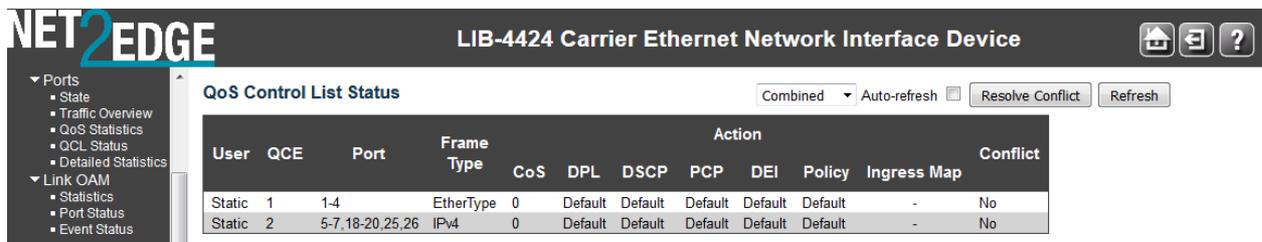
Auto-refresh: Check this checkbox to enable automatic refreshes of this page at 3 second intervals. See ‘

Detailed Port Statistics” on page 317 for more information on the ‘Detailed Port Statistics’ parameters.

Monitor > Ports > QCL Status

From the **Monitor > Ports > QCL Status** menu path you can view the QoS Control List Status table. This page shows the QCL status by different QCL users.

A QCE (QoS Control Entry) describes a QoS class associated with a particular QCE ID. There are six QCE frame types: Ethernet Type, VLAN, UDP/TCP Port, DSCP, TOS, and Tag Priority. Frames can be classified by one of four different QoS classes: "Low", "Normal", "Medium", and "High" for an individual application.



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 LIB-44xx. The displayed **QoS Control List Status** table parameters are explained below:

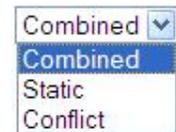
QoS Control List Status

From the dropdown list, select the QCL status to be displayed (*Combined*, *Static*, or *Conflict*), where:

Combined: Displays the 'Static' and 'Conflict' QCL status.

Static: Displays just the 'Static' QCL status (but not 'Conflict').

Conflict: Displays just the 'Conflict' QCL status (but not 'Static').



User

Indicates the QCL user type selected for display (e.g., *Combined*, *Static*, or *Conflict*).

QCE#

Indicates the index of the QCE on this line of the table (e.g., 1 or 2).

Frame Type

Indicates the type of frame to look for incoming frames. The QCE frame types are defined at the **Configuration > QoS > QoS Control List** menu path. Possible frame types are:

Any: The QCE will match all frame type.

Ethernet: Only Ethernet frames (with Ether Type 0x600-0xFFFF) are allowed.

LLC: Only LLC frames are allowed.

SNAP: Only SNAP frames are allowed. (The SubNetwork Access Protocol (SNAP) mechanism for multiplexing, on networks using IEEE 802.2 LLC.)

IPv4: The QCE will match only IPV4 frames.

IPv6: The QCE will match only IPV6 frames.

Port

Indicates the list of ports configured with the QCE (e.g., 1-3 or 2-4 on the LIB-4400).

Action

Indicates the classification action taken on ingress frames if parameters configured are matched with the frame's content. There are three action fields: Class, DPL and DSCP.

Class: Classified QoS Class; if a frame matches the QCE it will be put in the queue.

DPL: Drop Precedence level; if a frame matches the QCE then DP level will set to value displayed in the DP column.

DSCP: If a frame matches the QCE then DSCP will be classified with the value displayed under DSCP column.

Conflict

Displays QCE status. It may be 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**.

Note: to resolve a conflict, release the resource required by the QCE and click the 'Refresh' button.

Buttons

: Select the QCL status from this drop down list (**Combined**, **Static**, or **Conflict**).

Auto-refresh: Check this box to refresh the page automatically. Automatic refresh occurs at 3 second intervals.

Resolve Conflict: Release the resource required by the QCE and then click the '**Resolve Conflict**' button.

Refresh: Click to refresh the page; any changes made locally will be undone.

Monitor > Ports > Detailed Statistics

The **Monitor > Ports > Detailed Statistics** menu path displays the Detailed Port Statistics page for a Port. This page provides detailed traffic statistics for a specific LIB-44xx port. Use the port select box (e.g.,) to select which port details to display.

The LIB-44xx provides MAC bridging functionality per IEEE 802.1D. The LIB-44xx can operate in a VLAN unaware mode where it is an open bridge. The LIB-44xx can forward unicast, multicast, or broadcasts frames. RMON Counters based on frame type / frame size are maintained and reported. All error counters at the MAC layer are reported per ether-like MIB (RFC 2665) and/or IF-MIB (RFC 2863).

Receive Total		Transmit Total	
Rx Packets	0	Tx Packets	41716
Rx Octets	0	Tx Octets	2669824
Rx Unicast	0	Tx Unicast	0
Rx Multicast	0	Tx Multicast	41716
Rx Broadcast	0	Tx Broadcast	0
Rx Pause	0	Tx Pause	0
Receive Size Counters		Transmit Size Counters	
Rx 64 Bytes	0	Tx 64 Bytes	41716
Rx 65-127 Bytes	0	Tx 65-127 Bytes	0
Rx 128-255 Bytes	0	Tx 128-255 Bytes	0
Rx 256-511 Bytes	0	Tx 256-511 Bytes	0
Rx 512-1023 Bytes	0	Tx 512-1023 Bytes	0
Rx 1024-1526 Bytes	0	Tx 1024-1526 Bytes	0
Rx 1527- Bytes	0	Tx 1527- Bytes	0
Receive Queue Counters		Transmit Queue Counters	
Rx Q0	0	Tx Q0	41716
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	0
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	0		

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.

Receive Total and Transmit Total

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 Error Counters

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 frames received with valid CRC. 'Short frames' are frames that are smaller than 64 bytes.

Rx Oversize

The number of long frames received with valid CRC. 'Long frames' are frames that are longer than the configured maximum frame length for this port.

Rx Fragments

The number of short frames received with invalid CRC. 'Short frames' are frames that are smaller than 64 bytes.

Rx Jabber

The number of long frames received with invalid CRC. 'Long frames' are frames that are longer than the configured maximum frame length for this port.

Rx Filtered

The number of received frames filtered by the forwarding process.

Transmit Error Counters

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.

Buttons

: The port select box determines which port is affected (e.g., Port 1 - Port 4 on the LIB-4400).

Note: the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Refresh: Click to refresh the page immediately.

Clear: Clears the counters for the selected port.

Auto-refresh: Check this box to enable an automatic refresh of the page at 3 second intervals.

Monitor > Link OAM > Statistics

The **Monitor > Link OAM > Statistics** menu path displays the Detailed Link OAM Statistics for a Port.

This page provides detailed LOAM traffic statistics for a specific LIB-44xx port. Use the port select box to select which LIB-44xx 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.

Receive Total		Transmit Total	
Rx OAM Information PDU's	0	Tx OAM Information PDU's	0
Rx Unique Error Event Notification	0	Tx Unique Error Event Notification	0
Rx Duplicate Error Event Notification	0	Tx Duplicate Error Event Notification	0
Rx Loopback Control	0	Tx Loopback Control	0
Rx Variable Request	0	Tx Variable Request	0
Rx Variable Response	0	Tx Variable Response	0
Rx Org Specific PDU's	0	Tx Org Specific PDU's	0
Rx Unsupported Codes	0	Tx Unsupported Codes	0
Rx Link Fault PDU's	0	Tx Link Fault PDU's	0
Rx Dying Gasp	0	Tx Dying Gasp	0
Rx Critical Event PDU's	0	Tx Critical Event PDU's	0

The port Detailed Link OAM Statistics are explained below:

Receive Total and Transmit Total

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 OAMPDU's 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 OAMPDU's received and transmitted on this interface. Event Notification OAMPDU's 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 OAMPDU's 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 (Last Gasp). The LIB-44xx is equipped with the last gasp circuit for triggering a notification in the event of a power failure.

This will be useful for sending a notification. The uplink ports have highest priority to send the notifications of last gasp. The last gasp can be in the form of IEEE802.3 2008 Clause 57 Dying gasp event and/or an SNMP trap to NMS system. The management interface provides an option to choose the preferred mode of notification (either a SNMP trap or an IEEE 802.3 2008 clause 57 event).

Rx and Tx Critical Event PDU's

A count of the number of Critical event PDU's received and transmitted on this interface.

Buttons

: The port select box defines which port is affected by clicking the buttons (e.g., ports 1-14 on the).

Note: the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

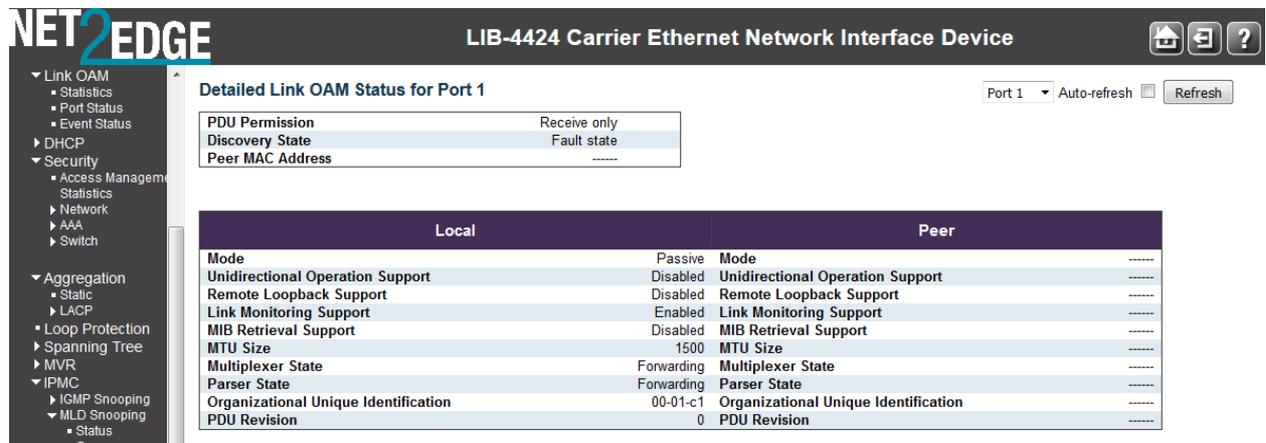
Auto-refresh: Check this checkbox to automatically refresh the page every three seconds.

Refresh: Click to refresh the page immediately.

Clear: Clears the counters for the selected port.

Monitor > Link OAM > Port Status

The **Monitor > Link OAM > Port Status** menu path displays the Detailed Link OAM Status for an LIB-44xx port. This page provides Link OAM configuration and operational status.



The displayed fields show the active configuration status for the selected port (e.g., 1-4 on the LIB-4400).

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.

Peer MAC Address

Displays the MAC address of the peer if known; displays "-----" if not known. Local and Peer

Mode

The Mode in which the Link OAM is operating, **Active** or **Passive**.

Unidirectional Operation Support

This feature is not user-configurable. 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 is 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 LIB-44xx is forwarding non-OAMPDUs to the lower sublayer. In the case of discarding, the LIB-44xx discards all the non-OAMPDU's.

Parser State

When in 'Forwarding' state, the LIB-44xx is forwarding non-OAMPDUs to a higher sublayer.
When in 'Loopback' state, the LIB-44xx is looping back non-OAMPDUs to the lower sublayer.
When in 'Discarding' state, the LIB-44xx is discarding non-OAMPDUs.

Organizational Unique Identification

This is the 24-bit Organizationally Unique Identifier (OUI) of the vendor (e.g., 00-C0-F2).

PDU Revision

Indicates the current revision of the Information TLV. The value of this field starts at zero and is 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).

Buttons

: Use the port select box to define which port is affected by clicking the buttons.

Note: the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Refresh: Click to refresh the page immediately.

Auto-refresh: Check this checkbox to enable an automatic refresh of the page at 3 second intervals.

Monitor > Link OAM > Event Status

The **Monitor > Link OAM > Event Status** menu path displays detailed Link OAM (LOAM) Link Status for an LIB-44xx port.

This page lets you view the current Link OAM Link Event configurations.

NET2EDGE LIB-4424 Carrier Ethernet Network Interface Device

Port 1 Auto-refresh 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
Symbol frame period errors	0	Symbol frame period errors	0
Symbol frame period error events	0	Symbol frame period error events	0
Local Event Seconds Summary Status		Remote Event Seconds Summary Status	
Event Seconds Summary Time Stamp	0	Event Seconds Summary Time Stamp	0
Event Seconds Summary Window	0	Event Seconds Summary Window	0
Event Seconds Summary Threshold	0	Event Seconds Summary Threshold	0
Event Seconds Summary Events	0	Event Seconds Summary Events	0
Event Seconds Summary Error Total	0	Event Seconds Summary Error Total	0
Event Seconds Summary Event Total	0	Event Seconds Summary Event Total	0

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.

Port dropdown

Port 1 Use port select box to select the port number (e.g., 1-4 for the LIB-4400).

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. The default value is one second. The lower bound is one second. 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. The default value is one frame error. The lower bound is zero frame errors. 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.

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.

Symbol frame period errors

This eight-octet field indicates the sum of symbol errors since the OAM sublayer was reset.

Symbol frame 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.

Event Seconds Summary Time Stamp

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.

Event Seconds Summary Window

This two-octet field indicates the duration of the period in terms of 100 ms intervals, encoded as a 16-bit unsigned integer.

Event Seconds Summary 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.

Event Seconds Summary Events

This two-octet field indicates the number of errored frame seconds in the period, encoded as a 16-bit unsigned integer.

Event Seconds Summary Error Total

This four-octet field indicates the sum of errored frame seconds that have been detected since the OAM sublayer was reset.

Event Seconds Summary Event Total

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 32-bit unsigned integer.

Buttons

: Use the port select box to define which port (e.g., ports 1-4 on the S4104) is affected by clicking the buttons.

Note: the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

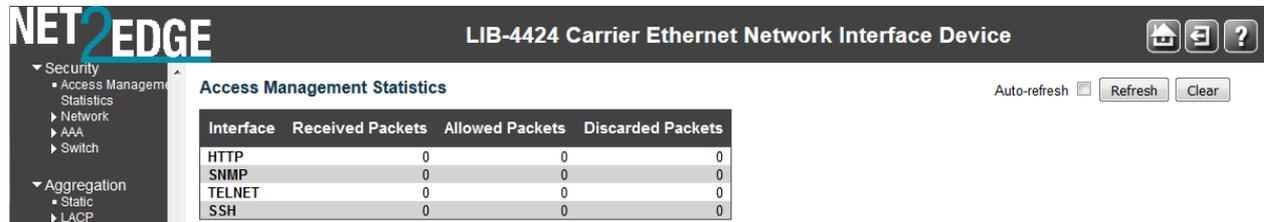
Refresh: Click to refresh the page.

Clear: Click to clear the data.

Monitor > Security > Access Management

The **Monitor > Security > Access Management** Statistics menu path displays the Access Management Statistics table. This page provides statistics on the various LIB-44xx access management interface methods.

The remote host can access the LIB-44xx via HTTP, HTTPS, SNMP, TELNET, and/or SSH.



The Access Management Statistics table parameters are explained below:

Interface

The interface type through which the remote host can access the LIB-44xx (**HTTP**, **HTTPS**, **SNMP**, **TELNET**, and/or **SSH**).

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.

Buttons

Auto-refresh: Check this box to enable an automatic refresh of the page at 3 second intervals.

Refresh: Click to refresh the page immediately.

Clear: Clear all statistics.

Monitor > Security > Network > Port Security

You can monitor the network device and ports' security from the **Monitor > Security > Network > Port Security** menu path.

Port Security > Switch

This page shows the current Port Security module and port status from the **Monitor > Security > Network > Port Security > Switch** menu path. Port Security is a module with no direct configuration. Configuration comes indirectly from other modules - the 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.

Port Security Switch Status

Auto-refresh Refresh

User Module Legend

User Module Name	Abbr
Limit Control	L
802.1X	8
Voice VLAN	V

Port Status

Port	Users	State	MAC Count	
			Current	Limit
1	L--	Ready	0	4
2	L--	Ready	0	4
3	L--	Ready	0	4
4	L--	Ready	0	4
5	L--	Ready	0	4
6	---	Disabled	-	-

The status page is divided into two sections - one with a legend of user modules and one with the actual port status.

User Module Legend

The legend shows all user modules that may request Port Security services.

User Module Name

The full name of a module that may request Port Security services.

Abbr

A one-letter abbreviation of the security user modules defined for the LIB-44xx ports.

Limit Control - L

802.1X - 8

DHCP Snooping - D

The abbreviation is used in the 'Users' column in the Port Status table.

Port Status

The table has one row for each LIB-44xx port and a number of columns. The columns are explained below:

Port

The port number for which the status applies. Click the port number to see the status for this particular port (e.g., ports 1-4 on the LIB-4400).

Note: the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

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**' description above) has enabled port security.

- L**- - - : Limit Control has enabled port security.
- S** - - - : 802.1X has enabled port security.
- D** - - - : DHCP Snooping has enabled port security.
- - - - : the corresponding user module is not enabled.

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 enabled by at least the Limit Control user module, and that module has indicated that the limit is reached and no more MAC addresses should be taken in.

Shutdown: The Port Security service is enabled by at least the Limit Control user module, and that module has indicated that the limit is exceeded. No MAC addresses can be learned on the port until it is administratively re-opened on the Limit Control configuration Web-page.

MAC Count (Current, Limit)

The two columns indicate the number of currently learned MAC addresses (forwarding as well as blocked) 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 the Limit Control user module is not enabled on the port, the Limit column will show a dash (-).

Buttons

Refresh: Click to refresh the page immediately.

Auto-refresh: Check this checkbox to enable an automatic refresh of the page at 3 second intervals.

In the Port Status section, when you click on a port in the table's Port column, the Port Security Port Status table displays for the specified LIB-44xx port (e.g., Port 1 below).

The screenshot shows the NET2EDGE web interface for a LIB-4424 Carrier Ethernet Network Interface Device. The main content area is titled 'Port Security Port Status Port 1'. There is a dropdown menu for 'Port 1', an 'Auto-refresh' checkbox, and a 'Refresh' button. Below this is a table with the following columns: MAC Address, VLAN ID, State, Time of Addition, and Age/Hold. The table currently contains one row with the text 'No MAC addresses attached'.

This page shows the MAC addresses secured by the Port Security module. Port Security is a module with no direct configuration. Configuration comes indirectly from other modules - the 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.

MAC Address

The MAC address and VLAN ID that is seen on this port. If no MAC addresses are learned, a single row stating "No MAC addresses attached" is displayed.

VLAN ID

The VLAN ID that is seen on this port (e.g., VID 1).

State

Indicates whether the corresponding MAC address is **Blocked** or **Forwarding**. In the blocked state, it will not be allowed to transmit or receive traffic.

Time of Addition

Shows the date and time when this MAC address was first seen on the port.

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 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.

Buttons

Port dropdown: Use the port select box () to select which port to show status for (e.g., ports 1-4 on the LIB-4400).

Note: the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Refresh: Click to refresh the page immediately.

Auto-refresh: Check this box to enable an automatic refresh of the page at 3 second intervals.

Monitor > Security > Network > NAS

You can monitor the network NAS device and ports' security status from the **Monitor > Security > Network > NAS** menu path.

Network Access Server system and port configuration is done at the **Configuration > Security > Network > NAS** menu path (see page 65).

NAS > Switch

This page provides an overview of the current NAS ports' states.

Port	Admin State	Port State	Last Source	Last ID	QoS Class	Port VLAN ID
1	Force Authorized	Link Down			-	
2	Port-based 802.1X	Link Down			-	
3	Single 802.1X	Link Down			-	
4	Multi 802.1X	Link Down			-	
5	Force Authorized	Link Down			-	
6	Force Authorized	Link Down			-	

The current NAS port states are explained below:

Port

Displays the LIB-44xx port number (linked). Click to display detailed NAS statistics for this port (e.g., ports 1-4 on the LIB-4400).

Note: the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Admin State

Displays the port's current administrative state. If NAS is globally enabled, this selection controls the port's authentication mode.

Force Authorized: In this mode, the switch will send one EAPOL Success frame when the port link comes up, and any client on the port will be allowed network access without authentication.

Force Unauthorized: In this mode, the switch will send one EAPOL Failure frame when the port link comes up, and any client on the port will be disallowed network access.

Port-based 802.1X: In 802.1X, the user is called the supplicant, the switch is the authenticator, and the RADIUS server is the authentication server. The authenticator acts as the man-in-the-middle, forwarding requests and responses between the supplicant and the authentication server. Frames sent between the supplicant and the switch are special 802.1X frames, known as EAPOL (EAP Over LANs) frames. EAPOL frames encapsulate EAP PDUs (RFC3748). Frames sent between the switch and the RADIUS server are RADIUS packets. RADIUS packets also encapsulate EAP PDUs together with other attributes like the switch's IP address, name, and the supplicant's port number on the switch.

Single 802.1X: In port-based 802.1X authentication, once a supplicant is successfully authenticated on a port, the whole port is opened for network traffic. This allows other clients connected to the port (for instance through a hub) to piggy-back on the successfully authenticated

client and get network access even though they really aren't authenticated. To overcome this security breach, use the Single 802.1X variant.

Single 802.1X is really not an IEEE standard, but features many of the same characteristics as does port-based 802.1X. In Single 802.1X, at most one supplicant can get authenticated on the port at a time. Normal EAPOL frames are used in the communication between the supplicant and the switch. If more than one supplicant is connected to a port, the one that comes first when the port's link comes up will be the first one considered. If that supplicant doesn't provide valid credentials within a certain amount of time, another supplicant will get a chance. Once a supplicant is successfully authenticated, only that supplicant will be allowed access. This is the most secure of all the supported modes. In this mode, the Port Security module is used to secure a supplicant's MAC address once successfully authenticated.

Multi 802.1X: In port-based 802.1X authentication, once a supplicant is successfully authenticated on a port, the whole port is opened for network traffic. This allows other clients connected to the port (for instance through a hub) to piggy-back on the successfully authenticated client and get network access even though they really aren't authenticated. To overcome this security breach, use the Multi 802.1X variant. Multi 802.1X is really not an IEEE standard, but features many of the same characteristics as does port-based 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" or "xx.xx.xx.xx.xx.xx" or "xxxxxxxxxxxx" (x is a hexadecimal digit). 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 maximum number of clients that can be attached to a port can be limited using the Port Security Limit Control functionality.

Port State

Displays the current state of the port. It can be one of the following values:

Globally Disabled: NAS is globally disabled at **Configuration > Security > Network > NAS >**

System Configuration > Mode = 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.

Last Source

Displays 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

Displays 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

Displays the QoS (Quality of Service) Class assigned to the port by the RADIUS server if enabled.

Port VLAN ID

Displays 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.

If the port is moved to the Guest VLAN, "(Guest)" is appended to the VLAN ID.

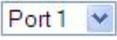
Buttons

Refresh: Click to refresh the page immediately.

Auto-refresh: Check this checkbox to automatic ally refresh the page every 3 seconds.

NAS > 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 the selected backend server (e.g., RADIUS Authentication Server) statistics only.

Use the port select box () to select which port details to be displayed (e.g., ports 1-4 on the LIB-4400).

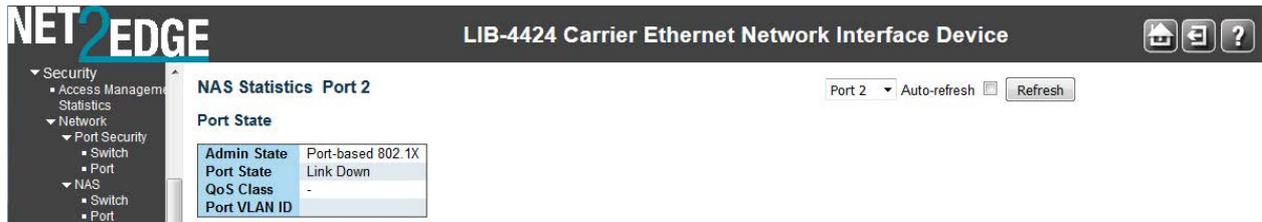
Note: the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Depending on the current port state, the **Monitor > Security > Network > NAS > Port** menu path displays Port State data, or Port State and Port Counters data. The port state can be 'Globally Disabled' 'Link Down', or 'Authorized' as shown on the screen examples below:

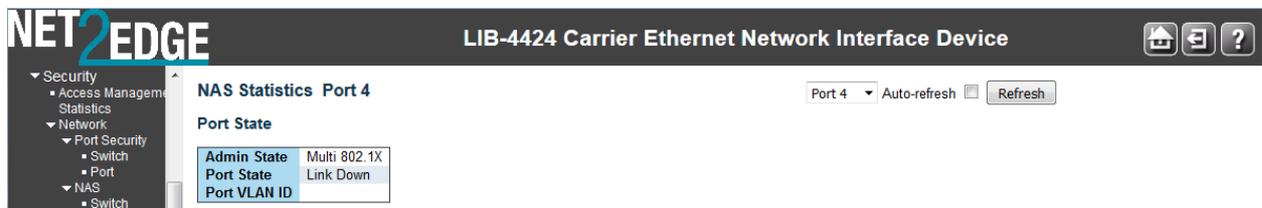
Port 1:



Port 2:



Port 4:



The NAS Port State and Port Counters are explained below:

Port State

Admin State

Displays the port's current administrative state. If NAS is globally enabled, this selection controls the port's authentication mode.

Force Authorized: In this mode, the switch will send one EAPOL Success frame when the port link comes up, and any client on the port will be allowed network access without authentication.

Force Unauthorized: In this mode, the switch will send one EAPOL Failure frame when the port link comes up, and any client on the port will be disallowed network access.

Port-based 802.1X: In 802.1X, the user is called the supplicant, the switch is the authenticator, and the RADIUS server is the authentication server. The authenticator acts as the man-in-the-middle, forwarding requests and responses between the supplicant and the authentication server. Frames sent between the supplicant and the switch are special 802.1X frames, known as EAPOL (EAP Over LANs) frames. EAPOL frames encapsulate EAP PDUs ([RFC3748](#)). Frames sent between the switch and the RADIUS server are **RADIUS** packets. RADIUS packets also encapsulate EAP PDUs together with other attributes like the switch's IP address, name, and the supplicant's port number on the switch.

Single 802.1X: In port-based 802.1X authentication, once a supplicant is successfully authenticated on a port, the whole port is opened for network traffic. This allows other clients connected to the port (for instance through a hub) to piggy-back on the successfully authenticated client and get network access even though they really aren't authenticated. To overcome this security breach, use the Single 802.1X variant.

Single 802.1X is really not an IEEE standard, but features many of the same characteristics as does port-based 802.1X. In Single 802.1X, at most one supplicant can get authenticated on the port at a time. Normal EAPOL frames are used in the communication between the supplicant and the switch. If more than one supplicant is connected to a port, the one that comes first when the port's link comes up will be the first one considered. If that supplicant doesn't provide valid credentials within a certain amount of time, another supplicant will get a chance. Once a supplicant is successfully authenticated, only that supplicant will be allowed access. This is the most secure of all the supported modes. In this mode, the Port Security module is used to secure a supplicant's MAC address once successfully authenticated.

Multi 802.1X: In port-based 802.1X authentication, once a supplicant is successfully authenticated on a port, the whole port is opened for network traffic. This allows other clients connected to the port (for instance through a hub) to piggy-back on the successfully authenticated client and get network access even though they really aren't authenticated. To overcome this security breach, use the Multi 802.1X variant. Multi 802.1X is really not an IEEE standard, but features many of the same characteristics as does port-based 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" or "xx.xx.xx.xx.xx.xx" or "xxxxxxxxxxxx" (x is a hexadecimal digit). 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 maximum number of clients that can be attached to a port can be limited using the Port Security Limit Control functionality.

Port State

Displays the current state of the port. It can be one of the following values:

Globally Disabled: NAS is globally disabled at **Configuration > Security > Network > NAS > System Configuration > Mode = 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.

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.

If the port is moved to the Guest VLAN, " (Guest) " is appended to the VLAN ID.

Port Counters

EAPOL Counters

These frame counters are available for these administrative states, as described in the table below: Force Authorized, Force Unauthorized, Port-based 802.1X, Single 802.1X, and Multi 802.1X.

Table 3. EAPOL Counters

Direction	Name	IEEE Name	Description
Rx	Total	dot1xAuthEapolFramesRx	The number of valid EAPOL frames of any type received by the LIB-44xx.
Rx	Response ID	dot1xAuthEapolRespldFramesRx	The number of valid EAPOL Response Identity frames received by the LIB-44xx.
Rx	Responses	dot1xAuthEapolRespFramesRx	The number of valid EAPOL response frames (other than Response Identity frames) received by the LIB-44xx.
Rx	Start	dot1xAuthEapolStartFramesRx	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 LIB-44xx.
Rx	Invalid Type	dot1xAuthInvalidEapolFramesRx	The number of EAPOL frames received by the LIB-44xx in which the frame type is not recognized.
Rx	Invalid Length	dot1xAuthEapLengthErrorFramesRx	The number of EAPOL frames

			received by the LIB-44xx in which the Packet Body Length field is invalid.
Tx	Total	dot1xAuthEapolFramesTx	The number of EAPOL frames of any type transmitted by the LIB-44xx.
Tx	Request ID	dot1xAuthEapolReqIdFramesTx	The number of EAPOL Request Identity frames transmitted by the LIB-44xx.
Tx	Requests	dot1xAuthEapolReqFramesTx	The number of valid EAPOL Request frames (other than Request Identity frames) transmitted by the LIB-44xx.

Backend Server Counters

These backend (RADIUS) frame counters are available for the administrative states (Port-based 802.1X, Single 802.1X, Multi 802.1X, and MAC-based Auth.) as described in the table below:

Table 4. Backend Server Counters

Direction	Name	IEEE Name	Description
Rx	Access Challenges	dot1xAuthBackendAccessChallenges	802.1X-based: Counts the number of times the LIB-44xx receives the first request from the backend server following the first response from the supplicant. Indicates that the backend server has communication with the LIB-44xx. 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	dot1xAuthBackendOtherRequestsTo Supplicant	802.1X-based: Counts the number of times the LIB-44xx sends an EAP Request packet following the first to the supplicant. Indicates 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 the LIB-44xx

			receives a success indication. Indicates the supplicant/client has successfully authenticated to the backend server.
Rx	Auth. Failures	dot1xAuthBackendAuthFails	802.1X- and MAC-based: Counts the number of times the LIB-44xx receives a failure message. Indicates the supplicant/client has not authenticated to the backend server.
Tx	Responses	dot1xAuthBackendResponses	802.1X-based: Counts the number of times the LIB-44xx attempts to send a supplicant's first response packet to the backend server. Indicates the LIB-44xx attempted communication with the backend server. Possible retransmissions are not counted. MAC-based: Counts all the backend server packets sent from the LIB-44xx 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

This field provides information about the last supplicant / client that attempted to authenticate (for the administrative states of Port-based 802.1X, Single 802.1X, Multi 802.1X, and MAC-based Auth.) as described in the table below:

Table 5. Last Supplicant/Client Information

Last Supplicant/Client Info		
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

Selected Counters

The Selected Counters table is visible when the port is in one of these administrative states:

- **Multi 802.1X**
- **MAC-based Auth.**

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 Addresses

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 displays *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 displays *No clients attached*.

VLAN ID

The VLAN ID that the corresponding client is currently secured through the Port Security module.

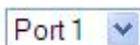
State

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).

Buttons



: Use the **Port select box** to select which port is affected when clicking the buttons (e.g., ports 1-4 on the LIB-4400).

Note: the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Auto Refresh: Check this box to enable an automatic refresh of the page at 3 second intervals.

Refresh: Click to refresh the page immediately.

Clear: This button is available in the following modes:

- Force Authorized
- Force Unauthorized
- Port-based 802.1X
- Single 802.1X

Click to clear the counters for the selected port.

Clear All: This button is available in the following modes:

- Multi 802.1X
- MAC-based Auth.X

Click to clear both the port counters and all of the attached client's counters. The "Last Client" will not be cleared, however.

Clear This: This button is available in the following modes:

- Multi 802.1X
- MAC-based Auth.X

Click to clear only the currently selected client's counters.

Monitor > Security > Network > ACL Status

You can display the current LIB-44xx ACL Status table from the **Monitor > Security > Network > ACL Status** menu path.

This page shows the ACL status by the various ACL users. (The related configuration is done at the **Configuration > Security > Network > ACL** menu path.)

The screenshot shows the NET2EDGE web interface for a LIB-4424 Carrier Ethernet Network Interface Device. The main content area displays the 'ACL Status' table. On the left is a navigation menu with categories like Security, Network, and AAA. On the right, there are controls for 'combined' view, 'Auto-refresh' checkbox, and a 'Refresh' button.

User	ACE	Frame Type	Action	Rate Limiter	Mirror	CPU	Counter	Conflict
ptp	1	EType	Deny	Disabled	Disabled	No	0	No
ptp	2	EType	Deny	Disabled	Disabled	No	0	No
ptp	3	EType	Deny	Disabled	Disabled	No	0	No
ptp	4	EType	Deny	Disabled	Disabled	No	0	No
ptp	5	EType	Deny	Disabled	Disabled	No	0	No
ptp	6	EType	Deny	Disabled	Disabled	No	0	No
ptp	7	EType	Deny	Disabled	Disabled	No	0	No
ptp	8	EType	Deny	Disabled	Disabled	No	0	No
static	1	IPv4/TCP 1234	Deny	10	Disabled	No	0	No

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 **256** on each LIB-44xx.

User

Indicates the ACL user type (e.g., **Static**, **MEP**, **PTP**, **ARP Inspection**, **IP Source Guard**).

Ingress Port

Indicates the ingress port of the ACE. Possible values are:

All: The ACE will match all ingress ports.

Port: The ACE will match a specific ingress port or set of ports (e.g., **1**, **3-4**).

Frame Type

Indicates the frame type of the ACE. The valid 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. An Etype frame may be followed by a suffix such as '-0x88f7' or '-0x8902'.

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 that are not ICMP/UDP/TCP (e.g., IPv4 SIP:192.168.1.210/32).

IPv6: The ACE will match all IPv6 standard frames.

Action

Indicates the forwarding action of the ACE. **Permit** forwards packets if all other ACL criteria are met. **Deny** drops packets if all other ACL criteria is met.

Permit: Frames matching the ACE may be forwarded and learned.

Deny: Frames matching the ACE are dropped.

Rate Limiter

Indicates the rate limiter number of the ACE. 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. 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".)

CPU

Forward packet that matched the specific ACE to CPU.

CPU Once

Forward first 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.

Buttons

Combined: Select the ACL status from this drop down list. The selections are Combined, Static, IP Source Guard, IPMC, MEP, ARP Inspection, PTP, DHCP, Loop Protect, (null), Link OAM, or Conflict. These selections are explained below:

Combined: displays the ACL status of all of the selections (if any exist). This is the default.

Static: displays the ACL status of just static ACL users.

IP Source Guard: displays the ACL status of just the IP Source Guard users.

IPMC: displays the ACL status of just IPMC ACL users.

MEP: displays the ACL status of just MEP ACL users.

ARP Inspection: displays the ACL status of just the ARP Inspection users.

PTP: displays the ACL status of just the PTP ACL users.

DHCP: displays the ACL status of just the DHCP users.

Loop Protect: displays the ACL status of just the loop protect mode users.

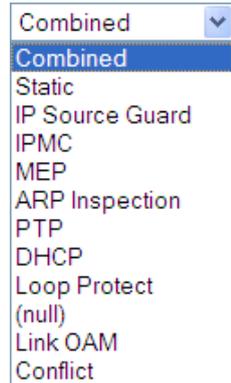
(null): displays the null set.

Link OAM: displays the ACL status of just the Link OAM type ACL users.

Conflict: displays the ACL status of just conflicted ACL users.

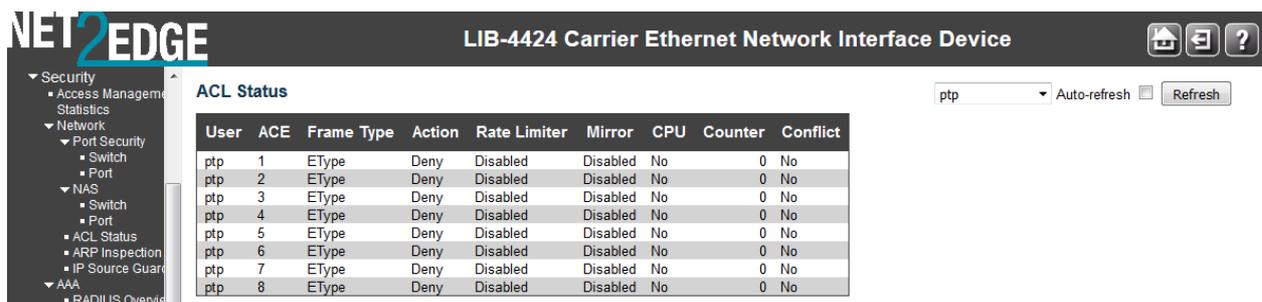
Auto-refresh: Check to refresh the page automatically. Automatic refresh occurs at 3 second intervals.

Refresh: Click to refresh the page; any changes made locally will be undone.

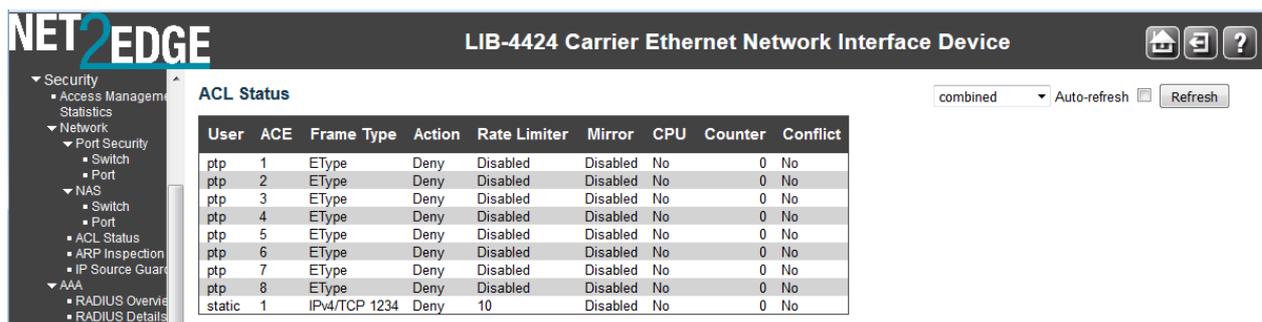


Example

The screen below shows monitoring for two configured ACLs.



The screen below shows monitoring for multiple configured ACLs.



See the **Configuration > Security > Network > ACL** menu path for the related configuration.

Monitor > Security > Network > DHCP

The **Monitor > Security > Network > DHCP** menu path provides the Relay Statistics pages.

DHCP > Relay Statistics

The **Monitor > Security > Network > DHCP > Relay Statistics** menu path provides client and server statistics for DHCP relay.

The screenshot shows the NET2EDGE web interface for a LIB-4424 Carrier Ethernet Network Interface Device. The main content area is titled "DHCP Relay Statistics" and includes an "Auto-refresh" checkbox, a "Refresh" button, and a "Clear" button. Below the title are two tables:

Server Statistics							
Transmit to Server	Transmit Error	Receive from Server	Receive Missing Agent Option	Receive Missing Circuit ID	Receive Missing Remote ID	Receive Bad Circuit ID	Receive Bad Remote ID
0	0	0	0	0	0	0	0

Client Statistics						
Transmit to Client	Transmit Error	Receive from Client	Receive Agent Option	Replace Agent Option	Keep Agent Option	Drop Agent Option
0	0	0	0	0	0	0

The page provides client and server Relay Statistics as explained below:

Server Statistics

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.

Client Statistics**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.

Buttons

Auto-refresh: Check this box to enable an automatic refresh of the page at 3 second intervals.

Refresh: Click to refresh the page immediately.

Clear: Clear all statistics.

Monitor > Security > Network > ARP Inspection

The **Monitor > Security > Network > ARP Inspection** menu path displays the Dynamic ARP Inspection Table.

Entries in the Dynamic ARP Inspection Table are shown on this page. The Dynamic ARP Inspection Table contains up to 1024 entries, and is sorted first by port, then by VLAN ID, then by MAC address, and then by IP address.

The screenshot shows the NET2EDGE web interface for a LIB-4424 Carrier Ethernet Network Interface Device. The main content area is titled "Dynamic ARP Inspection Table". It includes a navigation sidebar on the left with options like DHCP, Server, and Security. The main area contains a configuration section with the following fields: "Start from" (Port 1), "VLAN" (1), "MAC address" (00-00-00-00-00-00), and "IP address" (0.0.0.0), with a "with 20 entries per page" option. There are "Auto-refresh" checkboxes, a "Refresh" button, and navigation arrows. Below the configuration is a table with columns "Port", "VLAN ID", "MAC Address", and "IP Address". The table currently displays "No more entries".

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 let you select the starting point in the Dynamic ARP Inspection Table. Clicking the **Refresh** button will update the

displayed table starting from that or the closest next Dynamic ARP Inspection Table match. In addition, the two input fields will - upon a **Refresh** button click - assume the value of the first displayed entry, allowing for continuous refresh with the same start address.

Clicking 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" displays in the displayed table. Use the << button to start over.

The ARP Inspection table columns are explained below:

Port

Switch Port starting 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 in the format 00-00-00-00-00-00.

IP Address

User IP address of the entry in the format 0.0.0.0.

Buttons



: Use the port select dropdown box to select which port is affected by clicking the buttons (e.g., ports 1-4 on the LIB-4400).

Note: the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Auto-refresh: Check this checkbox to enable an automatic refresh of this page at 3 second intervals.

Refresh: Refreshes the displayed table starting from the input fields.

Clear: Flushes all dynamic entries.

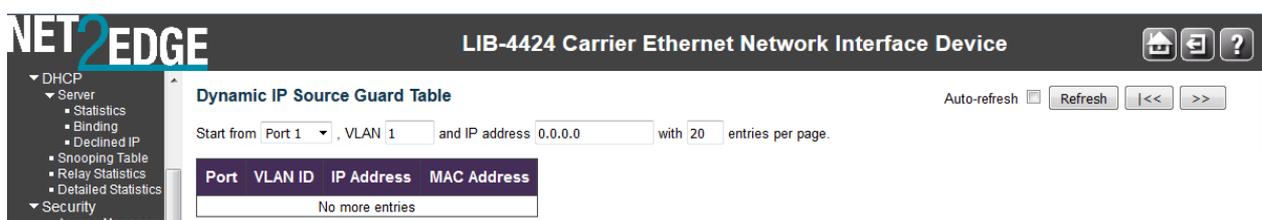
<<: Updates the table starting from the first entry in the Dynamic ARP Inspection Table.

>>: Updates the table, starting with the entry after the last entry currently displayed.

Monitor > Security > Network > IP Source Guard

The **Monitor > Security > Network > IP Source Guard** menu path displays the Dynamic ARP Inspection Table.

The Dynamic IP Source Guard Table is sorted first by port, then by VLAN ID, then by IP address, and then by MAC address.



Each page shows up to 99 entries from the Dynamic IP Source Guard 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 IP Source Guard Table.

The "Start from port address", "VLAN" and "IP address" input fields allow the user to select the starting point in the Dynamic IP Source Guard Table. Clicking the **Refresh** button will update the displayed table starting from that or the closest next Dynamic IP Source Guard Table match. In addition, the two input fields will - upon a **Refresh** button click - assume the value of the first displayed entry, allowing for continuous refresh with the same start address.

Clicking 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*" is shown in the displayed table. Use the << button to start over.

The IP Source Guard Table columns are explained below::

Port

Switch Port starting 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.

Buttons

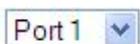
Auto-refresh: Check this checkbox to enable an automatic refresh of the page at 3 second intervals.

Refresh: Refreshes the displayed table starting from the input fields.

Clear: Flushes all dynamic entries.

<<: Updates the table starting from the first entry in the Dynamic IP Source Guard Table.

>>: Updates the table, starting with the entry after the last entry currently displayed.



: Use the port select dropdown box to select which port is affected by clicking the buttons.

Monitor > Security > AAA

The **Monitor > Security > AAA** menu path provides RADIUS Overview and RADIUS Details data.

AAA > RADIUS Overview

The RADIUS Authentication Overview page provides an overview of the status of the RADIUS servers configurable on the Authentication configuration page.

NET2EDGE LIB-4424 Carrier Ethernet Network Interface Device

RADIUS Server Status Overview Auto-refresh Refresh

#	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

RADIUS Authentication Servers

#

The RADIUS server number. Click to navigate to detailed statistics for this server.

IP Address

The IP address and UDP port number (in <IP Address>:<UDP Port> notation) of this server.

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.

RADIUS Accounting Servers

#

The RADIUS server number. Click to navigate to detailed statistics for this server.

IP Address

The IP address and UDP port number (in <IP Address>:<UDP Port> notation) of this server.

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 accounting attempts.

Dead (X seconds left): Accounting 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.

Buttons

Auto-refresh: Check this checkbox to enable an automatic refresh of the page at 3 second intervals.

Refresh: Click to refresh the page immediately.

AAA > RADIUS Details

The **Monitor > Security > AAA > RADIUS Details** menu path provides detailed RADIUS Authentication Statistics for a particular RADIUS server.

The statistics map closely to those specified in IETF [RFC4668 - RADIUS Authentication Client MIB](#). Use the server select box to switch between the backend servers to show details for.

RADIUS Authentication Statistics

Packet Counters

The RADIUS authentication server packet counters include seven receive counters and four transmit counters.

Direction	Name	RFC4668 Name	Description
Rx	Access Accepts	radiusAuthClientExtAccess Accepts	The number of RADIUS Access-Accept packets (valid or invalid) received from the server.
Rx	Access Rejects	radiusAuthClientExtAccess Rejects	The number of RADIUS Access-Reject packets (valid or invalid) received from the server.
Rx	Access Challenges	radiusAuthClientExtAccess Challenges	The number of RADIUS Access-Challenge packets (valid or invalid) received from the server.

Direction	Name	RFC4668 Name	Description
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.

Direction	Name	RFC4668 Name	Description
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.

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.

Direction	Name	RFC4670 Name	Description
Rx	Responses	radiusAccClientExtResponses	The number of RADIUS packets (valid or invalid) received from the server.

Direction	Name	RFC4670 Name	Description
Rx	Malformed Responses	radiusAccClientExtMalformed Responses	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	radiusAcctClientExtBad Authenticators	The number of RADIUS packets containing invalid authenticators received from the server.
Rx	Unknown Types	radiusAccClientExtUnknown Types	The number of RADIUS packets of unknown types that were received from the server on the accounting port.
Rx	Packets Dropped	radiusAccClientExtPackets Dropped	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	radiusAccClientExt Retransmissions	The number of RADIUS packets retransmitted to the RADIUS accounting server.
Tx	Pending Requests	radiusAccClientExtPending Requests	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.

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 accounting attempts. Dead (X seconds left): Accounting 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	radiusAccClientExtRoundTripTime	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.

Buttons

Server select box: The server select box determines which server is affected by clicking the buttons.

Auto-refresh: Check this box to enable an automatic refresh of the page at 3 second intervals.

Refresh: Click to refresh the page immediately.

Clear: Clears the counters for the selected server. The "Pending Requests" counter will not be cleared by this operation.

Monitor > Security > Switch > SNMP > RMON

This page lets you display RMON (Remote Monitoring) statistics, history, and alarms. You configure RMON at the **Configuration > Security > Switch > SNMP > RMON** menu path.

RMON > Statistics

This page provides an overview of RMON statistics entries (counters).

Each page shows up to 99 entries from the Statistics 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 Statistics table. The first displayed will be the one with the lowest ID found in the Statistics table.

The "Start from Control Index" allows the user to select the starting point in the Statistics table. Clicking the **Refresh** button will update the displayed table starting from that or the next closest Statistics table match. The >> button will use the last entry of the currently displayed entry as a basis for the next lookup. When the end is reached the text "No more entries" is shown in the displayed table.

Use the |<< button to start over.

ID	Source (ifIndex)	Drop	Octets	Pkts	Broad-cast	Multi-cast	CRC Errors	Under-size	Over-size	Frag.	Jabb.	Coll.	64 Bytes	65	128	256	512	1024
No more entries																		

The displayed RMON statistics counters are explained below:

ID

Indicates the index of the Statistics entry. Provides a link to the "Detailed RMON Statistics" display for the selected instance.

Data Source (ifIndex)

The data source which you want 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.

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 with a 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 from 65 to 127 octets in length.

128~255

The total number of packets (including bad packets) received that were from 128 to 255 octets in length.

256~511

The total number of packets (including bad packets) received that were from 256 to 511 octets in length.

512~1023

The total number of packets (including bad packets) received that were from 512 to 1023 octets in length.

1024~1588

The total number of packets (including bad packets) received that were from 1024 to 1588 octets in length.

Buttons

Auto-refresh: Check this box to enable an automatic refresh of the page at 3 second intervals.

Refresh: Click to refresh the page immediately.

|<<: Updates the table starting from the first entry in the Statistics table (i.e., the entry with the lowest ID).

>>: Updates the table, starting with the entry after the last entry currently displayed.

Detailed RMON Statistics

When you click an ID from the RMON Statistics Status Overview page, the “Detailed RMON Statistics” display for the selected instance.

NET2EDGE LIB-4424 Carrier Ethernet Network Interface Device

Detailed RMON Statistics ID 1

ID 1 Auto-refresh Refresh

Receive Total	
Port	1
Drops	0
Octets	10963696
Pkts	66663
Broadcast	920
Multicast	58422
CRC/Alignment	0
Undersize	0
Oversize	0
Fragments	2
Jabber	0
Collisions	0
64 Bytes	18528
65-127 Bytes	1137
128-255 Bytes	44429
256-511 Bytes	804
512-1023 Bytes	1363
1024-1518 Bytes	300

Buttons



ID 1: Use the ID select dropdown box to select which port's detailed RMON statistics display.

Auto-refresh: Check this box to enable an automatic refresh of the page at 3 second intervals.

Refresh: Click to refresh the page immediately.

|<<: Updates the table starting from the first entry in the Statistics table (i.e., the entry with the lowest ID).

>>: Updates the table, starting with the entry after the last entry currently displayed.

RMON > History

This page provides an overview of RMON history entries.

The RMON History Overview page provides an overview of RMON History entries. Each page shows up to 99 entries from the History 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 History table.

The first displayed will be the one with the lowest History Index and Sample Index found in the History table.

The **Start from History Index and Sample Index** lets you select the starting point in the History table. Clicking the **Refresh** button will update the displayed table starting from that or the next closest History table match.

The **>>** button will use the last entry of the currently displayed entry as a basis for the next lookup.

When the end is reached the text "*No more entries*" is shown in the displayed table.

Use the **|<<** button to start over.

RMON History Overview Auto-refresh

Start from Control Index and Sample Index with entries per page.

History Index	Sample Index	Sample Start	Drop	Octets	Pkts	Broad-cast	Multi-cast	CRC Errors	Under-size	Over-size	Frag.	Jabb.	Coll.	Utilization
1	1	6835163	0	6587088	101319	12	101033	0	0	0	0	0	0	0

The displayed fields are explained below:

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 total number of events in which packets were dropped by the probe due to lack of resources.

Drops

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.

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.

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 (.01% sampling interval).

Buttons

Refresh: Click to refresh the page immediately.

Auto-refresh: Check this box to enable an automatic refresh of the page at 3 second intervals.

|<<: Updates the table starting from the first entry in the History table (i.e., the entry with the lowest History Index and Sample Index).

>>: Updates the table, starting with the entry after the last entry currently displayed.

RMON > Alarm

This page provides an overview of RMON Alarm entries. Each page shows up to 99 entries from the Alarm table (the default is 20) as 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 Alarm table. The first displayed will be the one with the lowest ID found in the Alarm table.

The "Start from Control Index" allows the user to select the starting point in the Alarm table. Clicking the **Refresh** button will update the displayed table starting from that or the next closest Alarm table match.

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

Port Statistics Overview

Auto-refresh Refresh |<< >>

Start from Control Index with entries per page.

ID	Interval	Variable	Sample Type	Value	Startup Alarm	Rising Threshold	Rising Index	Falling Threshold	Falling Index
1	30	.13.6.1.2.1.2.2.1.10.10	Delta	72	RisingOrFalling	2	2	1	1

The displayed fields are explained below:

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, possible sample types are:

Rising Threshold

Rising threshold value.

Rising Index

Rising event index.

Falling Threshold

Falling threshold value.

Falling Index

Falling event index.

Buttons

Refresh: Click to refresh the page immediately.

Auto-refresh: Check this box to enable an automatic refresh of the page at 3 second intervals.

|<<: Updates the table starting from the first entry in the Alarm Table (i.e., the entry with the lowest ID).

>>: Updates the table, starting with the entry after the last entry currently displayed.

Example

The screenshot shows a web interface titled "Port Statistics Overview". At the top right, there is an "Auto-refresh" checkbox (unchecked), a "Refresh" button, and two navigation buttons: "|<<" and ">>". Below this, there is a text input field "Start from Control Index" with the value "0" and another text input field "with" with the value "20", followed by the text "entries per page".

ID	Interval	Variable	Sample Type	Value	Startup Alarm	Rising Threshold	Rising Index	Falling Threshold	Falling Index
1	30	.1.3.6.1.2.1.2.2.1.10.10	Delta	72	RisingOrFalling	2	2	1	1

Monitor > LACP > System Status

The LACP System Status table displays from the **Monitor > LACP > System Status** menu path. This page provides a status overview of all LACP instances. Configuration is done from the **Configuration > Aggregation > LACP** menu path.

The **Monitor > LACP** menu path displays the System Status, Port Status, and Port Statistics sub-menus. LACP (Link Aggregation Control Protocol) is an IEEE 802.3ad standard protocol that allows bundling several physical ports together to form a single logical port. The message “*No ports enabled or no existing partners*” displays if no status is available.

Aggr ID	Partner System ID	Partner Key	Partner Prio	Last Changed	Local Ports
No ports enabled or no existing partners					

The LACP System Status table parameters are explained below:

Aggr ID

The Aggregation ID associated with this aggregation instance. For LLAG the ID is shown as 'isid:aggr-id' and for GLAGs as 'aggr-id'.

Partner System ID

The system ID (MAC address) of the aggregation partner.

Partner Key

The Key that the partner has assigned to this aggregation ID.

Partner Prio

The priority 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 LIB-44xx.

Buttons

Refresh: Click to refresh this page immediately.

Auto-refresh: Check this checkbox to enable automatic refreshes of the page at 3 second intervals.

Monitor > LACP > Port Status

The LACP System Status table displays from the **Monitor > LACP > Port Status** menu path.

This page provides a status overview for LACP status for all LIB-44xx ports.

The screenshot shows the NET2EDGE web interface for a LIB-4424 Carrier Ethernet Network Interface Device. The left sidebar contains a navigation menu with categories like Aggregation, Loop Protection, Spanning Tree, MVR, and IPMC. The main content area is titled 'LACP Status' and features a table with the following data:

Port	LACP	Key	Aggr ID	Partner System ID	Partner Port	Partner Prio
1	No	-	-	-	-	-
2	No	-	-	-	-	-
3	No	-	-	-	-	-
4	No	-	-	-	-	-
5	No	-	-	-	-	-
6	No	-	-	-	-	-

Additional interface elements include an 'Auto-refresh' checkbox and a 'Refresh' button in the top right corner of the table area.

The LACP status table parameters are explained below:

Port

The LIB-44xx port number (e.g., 1-4 on the LIB-4400).

Note: the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

LACP

'**Yes**' means that LACP is enabled and the port link is up.

'**No**' means that LACP is not enabled or that the port link is down.

'**Backup**' means that the port could not join the aggregation group but will join if other port leaves. Meanwhile its LACP status is disabled.

Key

The key assigned to this port. Displays a number from 1-65535. Only ports with the same key can aggregate together.

Aggr ID

The Aggregation ID assigned to this aggregation group. IDs 1 and 2 are GLAGs while IDs 3-14 are LLAGs.

Partner System ID

The partner's System ID (MAC address).

Partner Port

The partner's port number connected to this port.

Partner Prio

The partner's port priority.

Buttons

Refresh: Click to refresh this page immediately.

Auto-refresh: Check this box to enable an automatic refresh of the page at 3 second intervals.

Monitor > LACP > Port Statistics

The LACP Statistics table displays from the **Monitor > LACP > Port Status** menu path.

This page provides an overview for LACP statistics for all ports.

The screenshot shows the NET2EDGE web interface for a LIB-4424 Carrier Ethernet Network Interface Device. The left sidebar contains a navigation menu with options like Aggregation, LACP, Loop Protection, Spanning Tree, MVR, IPMC, IGMP Snooping, and MLD Snooping. The main content area displays the 'LACP Statistics' table. Above the table are controls for 'Auto-refresh' (checkbox), 'Refresh', and 'Clear'.

Port	LACP Received	LACP Transmitted	Discarded	
			Unknown	Illegal
1	0	0	0	0
2	0	0	0	0
3	0	0	0	0
4	0	0	0	0
5	0	0	0	0
6	0	0	0	0

The LACP Statistics table parameters are explained below:

Port

The LIB-44xx port number (e.g., ports 1-4 on the LIB-4400).

Note: the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

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 and Illegal LACP frames have been discarded at each port.

Buttons

Auto-refresh: Check this box to enable an automatic refresh of this page at 3 second intervals.

Refresh: Click to refresh this page immediately.

Clear: Clears the counters for all ports.

Monitor > Loop Protection

The **Monitor > Loop Protection** menu path displays the loop protection port status of the LIB-44xx ports in the form of the Loop Protection Status table.

The screenshot shows the NET2EDGE web interface for a LIB-4424 Carrier Ethernet Network Interface Device. The main content area displays the 'Loop Protection Status' page. On the left, there is a navigation menu with options like Loop Protection, Spanning Tree, MVR, IPMC, IGMP Snooping, MLD Snooping, Status, Groups, Information, IPv6 SFM, Information, and LLDP. The main table shows the status of six ports. At the top right, there are buttons for 'Auto-refresh' (with a checkbox) and 'Refresh'.

Port	Action	Transmit	Loops	Status	Loop	Time of Last Loop
1	Shutdown+Log	Enabled	0	Down	-	-
2	Log Only	Enabled	0	Down	-	-
3	Shutdown+Log	Enabled	0	Down	-	-
4	Log Only	Enabled	0	Down	-	-
5	Shutdown	Enabled	0	Down	-	-
6	Shutdown	Enabled	0	Down	-	-

Loop protection port status parameters are explained below:

Port

The LIB-44xx port number of the logical port (e.g., ports 1-4 on the LIB-4400).

Note: the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Action

The currently configured port action taken on loop detection:

Shutdown Port: Shutdown the port.

Shutdown and Log : Shutdown the port and Log the event.

Log Only: Only Log the event.

Trap Only: Only send a trap.

Shutdown and Trap: Shutdown the port and Send trap.

Log and Trap: Send a Trap and Log the event.

All: Shutdown the port, send a trap, and Log the event.

Transmit

The currently configured port transmit mode (Enabled or Disabled).

Loops

The number of loops detected on this port.

Status

The current loop protection status of the port (**Up** or **Down**).

Loop

Whether a loop is currently detected on the port.

Time of Last Loop

The time that the last loop event was detected.

Buttons

Auto-refresh: Check this box to enable an automatic refresh of the page at 3 second intervals.

Refresh: Click to refresh the page immediately.

Monitor > Spanning Tree

The **Monitor > Spanning Tree** menu path displays the Bridge Status, Port Status, and Port Statistics sub-menus. Spanning tree protocols can include STP, MSTP, and RSTP.

Monitor > Spanning Tree > Bridge Status

The **Monitor > Spanning Tree > Bridge Status** menu path displays the STP Bridges table. This page provides a status overview of all STP bridge instances.

MSTI	Bridge ID	Root		Topology Flag	Topology Change Last
		ID	Port Cost		
CIST	32768.00-01-C1-00-FC-B0	32768.00-01-C1-00-FC-B0	- 0	Steady	-
MSTI1	32769.00-01-C1-00-FC-B0	32769.00-01-C1-00-FC-B0	- 0	Steady	-
MSTI2	32770.00-01-C1-00-FC-B0	32770.00-01-C1-00-FC-B0	- 0	Steady	-
MSTI3	32771.00-01-C1-00-FC-B0	32771.00-01-C1-00-FC-B0	- 0	Steady	-
MSTI4	32772.00-01-C1-00-FC-B0	32772.00-01-C1-00-FC-B0	- 0	Steady	-

The table displays a row for each STP bridge instance; the column information is explained below:

MSTI

The Bridge Instance. This is also a link to the 'STP Detailed Bridge Status'. MSTP allows formation of MST regions that can run multiple MST instances (MSTI).

Bridge ID

The Bridge ID of this Bridge instance (e.g., *80:00-00:C0:F2:21:B8:C4*).

Root ID

The Bridge ID of the currently elected root bridge.

Root Port

The LIB-44xx 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 (e.g., *Steady*).

Topology Change Last

The time since last Topology Change occurred.

Buttons

Refresh: Click to refresh the page immediately.

Auto-refresh: Check this box to enable an automatic refresh of the page at 3 second intervals.

Bridge Status Details

At **Monitor > Spanning Tree > Bridge Status** click on “CIST” in the MSTI column to display its details.

At **Monitor > Spanning Tree > Bridge Status** click on “MISTIX” in the MSTI column to display its details.

The screenshot shows the NET2EDGE web interface for a LIB-4424 Carrier Ethernet Network Interface Device. The main content area displays the 'STP Detailed Bridge Status' for the CIST instance. The 'STP Bridge Status' table contains the following data:

STP Bridge Status	
Bridge Instance	CIST
Bridge ID	32768.00-01-C1-00-FC-B0
Root ID	32768.00-01-C1-00-FC-B0
Root Cost	0
Root Port	-
Regional Root	32768.00-01-C1-00-FC-B0
Internal Root Cost	0
Topology Flag	Steady
Topology Change Count	0
Topology Change Last	-

Below this table is the 'CIST Ports & Aggregations State' table, which currently shows 'No ports or aggregations active'.

The STP Detailed Bridge Status table parameters are shown below with sample parameters.

Bridge Instance

The Bridge instance (e.g., **MSTI1** or **CIST**).

Bridge ID

The Bridge ID of this Bridge instance (e.g., **80:01-00:C0:F2:00:00:01**).

Root ID

The Bridge ID of the currently elected root bridge (e.g., **80:01-00:C0:F2:00:00:01**).

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 (e.g., **0**).

Root Port

The switch port currently assigned the root port role (e.g., **-** indicating none reported).

Regional Root

The Bridge ID of the currently elected regional root bridge, inside the MSTP region of this bridge (e.g., **80:00-00:C0:F2:00:00:01**) (displays for CIST 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 (e.g., **0**) (displays for CIST only).

Topology Flag

The current state of the Topology Change Flag of this Bridge instance (e.g., **Steady**).

Topology Change Count

The number of times where the topology change flag has been set (during a one-second interval) (e.g., **0**).

Topology Change Last

The time passed since the Topology Flag was last set (e.g., **0d 01:15:55**, or **-** indicating none encountered).

The **CIST Ports & Aggregations State** table parameters are shown below with sample parameters.

CIST Ports & Aggregations State

Port	Port ID	Role	State	Path Cost	Edge	Point-to-Point	Uptime
1	128:001	DesignatedPort	Forwarding	2000	Yes	Yes	0d 00:16:41
2	128:002	BackupPort	Discarding	2000	No	Yes	0d 00:16:41

Port

The switch port number of the logical STP port (e.g., **1**).

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 (e.g., **128:009**).

Role

The current STP port role. The port role can be **AlternatePort**, **BackupPort**, **RootPort**, or **DesignatedPort**.

State

The current STP port state. The port state can be **Discarding**, **Learning**, or **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 (e.g., **200000**).

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 (e.g., **Yes** or **No**).

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. (e.g., **Yes** or **No**)

Uptime

The time since the bridge port was last initialized (e.g., **0d 03:54:57**).

If nothing is configured, the message “No ports or aggregations active” displays.

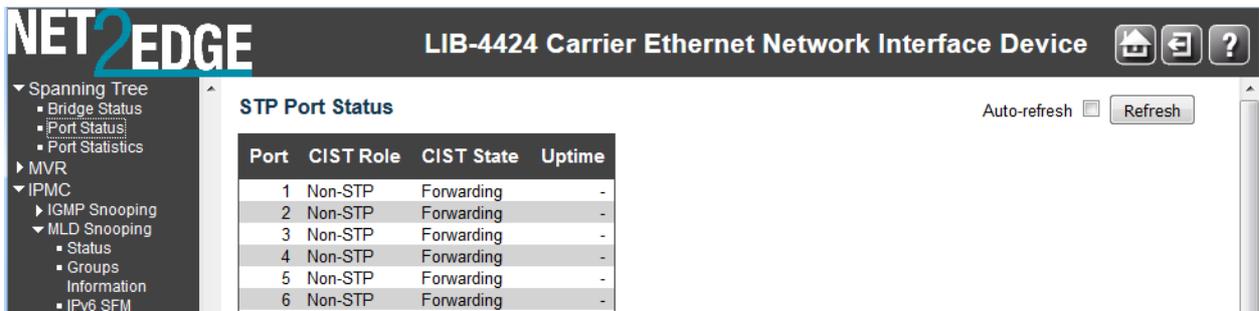
Monitor > Spanning Tree > Port Status

The **Monitor > Spanning Tree > Port Status** menu path displays the STP Port Status table, which provides the STP CIST port status for the LIB-44xx physical ports.

An MSTn instance is local to a region. ISTs in different regions are interconnected via a Common Spanning Tree (CST). The Common and Internal Spanning Tree (CIST) includes the collection of ISTs in each MST region, and the CST that connects the ISTs.

As a result of the spanning-tree calculation, ports will assume various roles in the topology. A root port is a port facing towards the root that is connected to the best path back to the root. ‘Best path’ means

- 1) the path with the lowest cost back to the root, 2) the path going through the device with the lowest BID if there is more than device advertising the lowest cost, and 3) the lowest port ID on that device if there is more than one connection to the device.



The STP Port Status table parameters are explained below:

Port

The LIB-44xx port number of the logical STP port (e.g., ports 1-4 on the LIB-4400).

Note: the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

CIST Role

The current STP port role of the CIST port. The port’s CIST Role value can be:

AlternatePort: An alternative path to the root bridge. This path is different than using the root port.

BackupPort: A backup/redundant path to a segment where another bridge port already connects.

RootPort: Port by which frames leave a device to reach the root (forwarding port). This refers to a forwarding port that is the best port from Nonroot-bridge to Rootbridge.

DesignatedPort: Port by which frames enter a device to reach the root (forwarding port).

A forwarding port for every LAN segment. A Non-Designated port is a Port blocking frames to prevent a loop in the topology (blocking port).

Disabled: Not strictly part of STP, a network administrator can manually disable a port.

Non-STP: STP is disabled for this port.

CIST State

The current STP port state of the CIST port. The port state can be one of the following values:

Blocking, Learning, Forwarding, Discarding, Listening, or Disabled. STP per IEEE 802.1d proceeds through various states to establish the topology.

Blocking: A port that would cause a switching loop; no user data is sent or received but it may go into forwarding mode if the other links in use were to fail and the spanning tree algorithm determines the port may transition to the forwarding state. BPDU data is still received in blocking state.

Ports start out in the Blocking state and then proceed to the listening state where they exchange BPDUs with other devices to establish the root switch, root ports, and designated ports. Blocking state does not mean the port is shut down; it just means that no frames are allowed to be sent or received via that port. Spanning tree information continues to be received via that port from the designated bridge on that segment, which allows STP to communicate a change in topology should one occur.

Learning: After the Blocking state, ports then enter a Learning phase where they listen to frames reaching their ports to build the MAC tables on the device. While the port does not yet forward frames (packets) it does learn source addresses from frames received and adds them to the filtering database (switching database).

Forwarding: After the Learning state, the root and designated ports enter the Forwarding state, while the non-designated ports are set to blocking state. A port receiving and sending data, normal operation. STP still monitors incoming BPDUs that would indicate it should return to the blocking state to prevent a loop.

Discarding: The device is discarding all non-OAMPDUs.

Listening: The switch processes BPDUs and awaits possible new information that would cause it to return to the blocking state.

Disabled - Not strictly part of STP, a network administrator can manually disable a port.

Uptime

The time since the bridge port was last initialized (e.g., **9d 04:09:48**, or 9 days, 4 hours, 9 minutes and 48 seconds).

Buttons

Refresh: Click to refresh the page immediately.

Auto-refresh: Check this checkbox to enable automatic page refreshes at 3 second intervals.

Monitor > Spanning Tree > Port Statistics

The **Monitor > Spanning Tree > Port Statistics** menu path displays the STP Port Status table, which provides the STP port statistics counters of LIB-44xx bridge ports. STP statistics are provided for STP, MSTP, and RSTP.

The screenshot shows the NET2EDGE web interface for a LIB-4424 Carrier Ethernet Network Interface Device. The left sidebar contains a navigation menu with 'Spanning Tree' expanded to show 'Port Statistics'. The main content area is titled 'STP Statistics' and includes an 'Auto-refresh' checkbox, 'Refresh', and 'Clear' buttons. Below this is a table with the following structure:

Port	Transmitted				Received				Discarded	
	MSTP	RSTP	STP	TCN	MSTP	RSTP	STP	TCN	Unknown	Illegal
No ports enabled										

The STP port Statistics counters are explained below:

Port

The LIB-44xx port number of the logical STP port (e.g., ports 1-4 on the LIB-4400).

Note: the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

MSTP

The number of MSTP Configuration BPDUs received/transmitted on the port. The Multiple Spanning Tree Protocol (MSTP) allows formation of MST regions that can run multiple MST instances (MSTI). Multiple regions and other STP bridges are interconnected using one single common spanning tree (CST).

RSTP

The number of RSTP Configuration BPDUs received/transmitted on the port. IEEE document 802.1w introduced RSTP as an evolution of STP. The Rapid Spanning Tree Protocol, which provides for faster spanning tree convergence after a topology change. Standard IEEE 802.1D-2004 now incorporates RSTP and obsoletes STP, while at the same time remains backwards compatible with STP.

STP

The number of legacy STP Configuration BPDUs received/transmitted on the port. The Spanning Tree Protocol (STP, per IEEE 802.1D) creates a spanning tree within a mesh network of connected layer-2 bridges (typically Ethernet switches), and disables those links that are not part of the spanning tree, leaving a single active path between any two network nodes.

TCN

The number of (legacy) Topology Change Notification (TCN) BPDUs received/transmitted on the port. TCN BPDUs are used to inform other devices of port changes. TCNs are injected into the network by a non-root switch and propagated to the root. On receipt of the TCN, the root switch will set a Topology Change flag in its normal BPDUs. This flag is propagated to all other switches to instruct them to rapidly age out their forwarding table entries.

Discarded Unknown

The number of unknown Spanning Tree BPDUs received (and discarded) on the port.

Discarded Illegal

The number of illegal Spanning Tree BPDUs received (and discarded) on the port.

Buttons

Refresh: Click to refresh the page immediately.

Auto-refresh: Check this checkbox to enable automatic page refreshes at 3 second intervals.

Clear: Click to reset the counters.

Example

Port	Transmitted				Received				Discarded	
	MSTP	RSTP	STP	TCN	MSTP	RSTP	STP	TCN	Unknown	Illegal
1	862	0	0	0	3	0	0	0	0	0
2	3	0	0	0	862	0	0	0	0	0

Monitor > MVR

You can view Statistics, MVR Channel Groups, and MVR SFM Information from the **Monitor > MVR** menu path. MVR is configured from the **Configuration > MVR** menu path.

Statistics

This page provides MVR Statistics information from the **Monitor > MVR > Statistics** menu path.

VLAN ID	IGMP/MLD		IGMPv1	IGMPv2/MLDv1		IGMPv3/MLDv2	IGMPv2/MLDv1
	Queries Received	Queries Transmitted		Reports Received	Reports Received		
2	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

The MVR Statistics information is explained below:

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 Joins.

IGMPv2/MLDv1 Reports Received

The number of Received IGMPv2 Joins and MLDv1 Reports, respectively.

IGMPv3/MLDv2 Reports Received

The number of Received IGMPv1 Joins and MLDv2 Reports, respectively.

IGMPv2/MLDv1 Leaves Received

The number of Received IGMPv2 Leaves and MLDv1 Dones, respectively.

Buttons

Auto-refresh: Automatic refresh occurs every 3 seconds.

Refresh: Click to refresh the page immediately.

Clear: Clears all Statistics counters.

MVR Channel Groups

This page provides MVR Channel (Groups) information from the **Monitor > MVR > MVR Channel Groups** menu path.

Entries in the MVR Channels (Groups) Information Table are shown on this page. The MVR Channels (Groups) Information Table is sorted first by VLAN ID, and then by group.

Each page shows up to 99 entries from the MVR Group 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 MVR Channels (Groups) Information Table.

The "Start from VLAN", and "Group Address" input fields let you select the starting point in the MVR Channels (Groups) Information table. Clicking the **Refresh** button will update the displayed table starting from that or the closest next MVR Channels (Groups) Information table match. In addition, the two input fields will - upon a >> button click - assume the value of the first displayed entry, allowing for continuous refresh with the same start address.

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" displays in the table. Use the << button to start over.

The MVR Channel (Groups) information is explained below:

Group Address

The IPv4 / IPv6 address of the group (e.g., IPv4 multicast addresses range 224.0.0.0 to 239.255.255.255).

VLAN ID

The VLAN ID of the group.

Groups

The Group ID of the group displayed.

Port Members

Ports under this group (e.g., ports 1-6 on the LIB-4400).

Note: the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Buttons

Auto-refresh: Automatic refresh occurs every 3 seconds.

Refresh: Refreshes the displayed table starting from the input fields.

|<<: Updates the table starting from the first entry in the MVR Channels (Groups) Information Table.

>>: Updates the table, starting with the entry after the last entry currently displayed.

MVR SFM Information

This page provides MVR Channel (Groups) information from the **Monitor > MVR > MVR SFM Information** menu path.

The MVR SFM (Source-Filtered Multicast) Information Table also contains the SSM (Source-Specific Multicast) information. This table is sorted first by VLAN ID, then by group, and then by Port.

Different source addresses belong to the same group are treated as single entry.

Each page shows up to 99 entries from the MVR SFM Information Table (default is 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 MVR SFM Information Table.

The "Start from VLAN", and "Group Address" input fields allow the user to select the starting point in the MVR SFM Information Table. Clicking the **Refresh** button will update the displayed table starting from that or the closest next MVR SFM Information Table match. In addition, the two input fields will - upon a **Refresh** button click - assume the value of the first displayed entry, allowing for continuous refresh with the same start address.

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" is shown in the displayed table. Use the |<< button to start over.

The MVR SFM (Source-Filtered Multicast) Information table entries are explained below:

Group Address

The IPv4 / IPv6 address of the group (e.g., IPv4 multicast addresses range 224.0.0.0 to 239.255.255.255).

VLAN ID

VLAN ID of the group.

Group

Group address of the group displayed.

Port

The LIB-44xx port number (e.g., ports 1-4 on the LIB-4400).

Note: the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Mode

Indicates the filtering mode maintained per (VLAN ID, Port number, Group Address) basis. Mode can be set to either **Include** or **Exclude**.

Source Address

The IP Address of the source. Currently, the LIB-44xx limits the total number of IP source addresses for filtering to be **128**. When there is no any source filtering address, the text "*None*" displays in the Source Address field.

Type

Indicates the type of MVR performed. It can be either **Allow** or **Deny**.

Hardware Filter/Switch

Indicates whether the data plane destined to the specific group address from the source IPv6 address could be handled by chip.

Buttons

Auto-refresh: Check the checkbox to cause an automatic refresh to occur every 3 seconds.

Refresh: Refreshes the displayed table starting from the input fields.

|<<: Updates the table starting from the first entry in the MVR SFM Information table.

>>: Updates the table, starting with the entry after the last entry currently displayed.

Monitor > IPMC > IGMP Snooping

The **Monitor > IPMC > IGMP Snooping** menu path provides the Status, Groups Information, and IPv4 SSM Information sub-menus.

The IGMP (Internet Group Management Protocol) communications protocol is used to manage the membership of Internet Protocol multicast groups. IGMP is used by IP hosts and adjacent multicast routers to establish multicast group memberships. It is an integral part of the IP multicast specification, like ICMP for unicast connections. IGMP allows more efficient use of resources when supporting uses such as online video.

IGMP Snooping Status

This page provides IGMP Snooping status in terms of statistics and router port status.

The screenshot shows the NET2EDGE web interface for a LIB-4424 Carrier Ethernet Network Interface Device. The main content area displays the 'IGMP Snooping Status' page. On the left is a navigation menu with categories like IGMP Snooping, MLD Snooping, LLDP, Ethernet Services, Performance Monitor, PTP, and MAC Table. The 'IGMP Snooping Status' section has a sub-section for 'Statistics' with a table showing details for VLANs 10, 20, 30, and 40. Below this is a 'Router Port' table showing ports 1 through 5.

VLAN ID	Querier Version	Host Version	Querier Status	Queries Transmitted	Queries Received	V1 Reports Received	V2 Reports Received	V3 Reports Received	V2 Leaves Received
10	v3	v3	ACTIVE	0	0	0	0	0	0
20	v1	v1	ACTIVE	0	0	0	0	0	0
30	v2	v2	ACTIVE	0	0	0	0	0	0
40	v3	v3	ACTIVE	0	0	0	0	0	0

Port	Status
1	-
2	-
3	-
4	-
5	-

The IPMC Snooping status **Statistics** information is explained below:

VLAN ID

The VLAN ID (VID) of the entry (e.g., VLAN ID 1).

Querier Version

The working Querier Version currently in use (e.g., **v3** shown above). In order for IGMP, and IGMP snooping, to function, a multicast router must exist on the network and generate IGMP queries. The tables created for snooping that (hold the member ports for a multicast group) are associated with the querier. Without a querier, the tables are not created and snooping does not work. IGMP general queries must be unconditionally forwarded by all switches involved in IGMP snooping.

Host Version

The working Host Version currently in use (e.g., **v3** shown above).

Three versions of IGMP exist - versions **v1**, **v2**, and **v3**. One difference between the versions is how a host node signals that it no longer wants to be a member of a multicast group.

In **v1**, the host node stops sending reports. If a router does not receive a report from a host node after a predefined length of time (time-out value) it assumes that the host node no longer wants to receive multicast frames and removes it from the membership list of the multicast group.

In **v2**, a host node exits from a multicast group by sending a leave request. After receiving a leave request from a host node, the router removes the node from appropriate membership list. The router also stops sending multicast packets from the port if it determines there are no further host nodes on the port.

v3 adds the ability of host nodes to "join" or "leave" specific sources in a multicast group. IGMP support in Microsoft Windows hosts includes:

IGMPv1 in Windows 95, Windows NT 4.0 (SP3 and earlier).

IGMPv2 in Windows 98, Windows ME, Windows NT 4.0 (SP4 and later), Windows 2000.

IGMPv3 in Windows XP, Windows Server 2003, Windows Vista.

For more information see <http://support.microsoft.com/>.

Querier Status

Displays the Querier status as "**ACTIVE**" or "**IDLE**". Displays "**DISABLE**" to note that the specific interface is administratively disabled.

Queries Transmitted

The number of Transmitted Queries (e.g., 5 shown above).

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.

The IPMC Snooping status table **Router Port** information is explained below:

Port

The LIB-44xx port number (e.g., ports 1-4 on the LIB-4400).

Note: the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Status

Indicates whether a specific port is a router port. The IGMP Snooping status Router Port table displays 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. The status reported can be:

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.

-- indicates there is no status to display.

Buttons

Refresh: Click to refresh the page immediately.

Clear: Clears all Statistics counters.

Auto-refresh: Check this checkbox to enable automatic page refreshes at 3 second intervals.

IGMP Snooping > Groups Information

Entries in the IGMP Group Table are displayed on this page. The IGMP Group Table is sorted first by VLAN ID, and then by Group.

The screenshot shows the NET2EDGE web interface for a LIB-4424 Carrier Ethernet Network Interface Device. The main content area is titled "IGMP Snooping Group Information". It features a search bar with "Start from VLAN" set to 1 and "group address" set to 224.0.0.0, with "20" entries per page. Below the search bar is a table titled "Port Members" with columns for "VLAN ID" and "Groups". The table currently displays "No more entries". Navigation buttons include "Auto-refresh" (unchecked), "Refresh", "I<<", and ">>".

Each page shows up to 99 entries from the IGMP Group table selected through the "entries per page" input field (the default is 20). When first visited, the web page will show the first 20 entries from the beginning of the IGMP Group Table.

The "Start from VLAN", and "group address" input fields let you select the starting point in the IGMP Group Table. Clicking the **Refresh** button will update the displayed table starting from that or the closest next IGMP Group Table match. In addition, the two input fields will - upon a **Refresh** button click - assume the value of the first displayed entry, allowing for continuous refresh with the same start address.

The 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" is shown in the displayed table. Use the |<< button to start over.

VLAN ID

VLAN ID of the IGMP group.

Groups

Group address of the IGMP group displayed.

Port Members

Ports under this IGMP group (e.g., ports 1-4 on the LIB-4400).

Note: the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Buttons

Auto-refresh: Check this checkbox to enable an automatic refresh of the page at 3 second intervals.

Refresh: Refreshes the displayed table starting from the input fields.

|<<: Updates the table, starting with the first entry in the IGMP Group Table.

>>: Updates the table, starting with the entry after the last entry currently displayed.

IGMP Snooping IPv4 SSM Information

Entries in the IGMP SSM Information Table are displayed on this page. The IGMP SSM Information Table is sorted first by VLAN ID, then by group, and then by Port Number. Different source addresses that belong to the same group are treated as single entry.

Each page shows up to 99 entries from the IGMP SSM (Source Specific Multicast) Information 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 IGMP SSM Information Table.

The "Start from VLAN", and "Group" input fields allow the user to select the starting point in the IGMP SSM Information Table. Clicking **Refresh** the button will update the displayed table starting from that or the closest next IGMP SSM Information Table match. In addition, the two input fields will - upon a **Refresh** button click - assume the value of the first displayed entry, allowing for continuous refresh with the same start address.

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*" displays in the displayed table. Use the << button to start over.

VLAN ID

VLAN ID of the group (e.g., VLAN ID 1).

Group

Group address of the group displayed (e.g., 239.255.255.250).

Port

LIB-44xx port number (e.g., 1-4).

Mode

Indicates the filtering mode maintained per (VLAN ID, port number, Group Address) basis. It can be either '**Include**' or '**Exclude**'.

Source Address

The IP Address of the source. The total number of IP source addresses for filtering is limited to **128**.

Type

Indicates the Type. It can be either '**Allow**' or '**Deny**'.

Hardware Filter/Switch

Indicates whether the data plane destined to the specific group address from the source IPv4 address could be handled by chip.

Buttons

Auto-refresh: Check this box to enable an automatic refresh of the page at 3 second intervals.

Refresh: Refreshes the displayed table starting from the input fields.

<<: Updates the table starting from the first entry in the IGMP SSM Information Table.

>>: Updates the table, starting with the entry after the last entry currently displayed.

Monitor > IPMC > MLD Snooping

The **Monitor > IPMC > MLD Snooping** menu path provides the Status, Groups Information, and IPv6 SSM Information sub-menus. MLD (Multicast Listener Discovery) for IPv6 is used by IPv6 routers to discover multicast listeners on a directly-attached link (much as IGMP is used in IPv4). The MLD protocol is embedded in ICMPv6 instead of using a separate protocol.

MLD Snooping > Status

The **MLD Snooping > Status** page provides MLD Snooping status.

The screenshot shows the NET2EDGE web interface for device LIB-4424. The main content area displays the 'MLD Snooping Status' page. On the left is a navigation menu with options like MLD Snooping, LLDP, Ethernet Services, Performance, Monitor, PTP, MAC Table, VLANs, MVRP, and sFlow. The 'MLD Snooping Status' page has a title bar with 'LIB-4424 Carrier Ethernet Network Interface Device' and navigation icons. Below the title bar are 'Auto-refresh' (checkbox), 'Refresh', and 'Clear' buttons. The 'Statistics' table shows two entries for VLANs 10 and 11, both with Querier Version v2, Host Version v2, and Status ACTIVE. The 'Router Port' table shows ports 1 through 5, all with a Status of '-'.

VLAN ID	Querier Version	Host Version	Querier Status	Queries Transmitted	Queries Received	V1 Reports Received	V2 Reports Received	V1 Leaves Received
10	v2	v2	ACTIVE	0	0	0	0	0
11	v2	v2	ACTIVE	0	0	0	0	0

Port	Status
1	-
2	-
3	-
4	-
5	-

The MLD Snooping Status table **Statistics** parameters are explained below:

VLAN ID

The VLAN ID of the entry (e.g., VLAN ID 1).

Querier Version

The working Querier Version currently (e.g., **v2**).

Host Version

The working Host Version currently in use (e.g., **v2** shown above).

MLD **v1** was the original release of MLD as an asymmetric protocol, specifying different behaviours for multicast listeners and for routers per IETF [RFC 2710](#).

MLD**v2** is designed to be interoperable with MLDv1. MLDv2 adds the ability for a node to report interest in listening to packets with a particular multicast address only from specific source addresses or from all sources except for specific source addresses. Refer to IETF [RFC 3810](#).

Windows support includes:

MLDv1 in Windows 98, Windows ME, Windows NT 4.0 (SP4 and later), Windows 2000.

MLDv2 in Windows XP, Windows Server 2003, Windows Vista.

Windows XP supports the host side of MLDv1 and can function as a multicast source or receiver. Multicast routing support is not present on XP.

IPv6 in Windows Server 2008 and Windows Vista supports both MLD and MLDv2.

IPv6 in Windows Server 2008 and Windows Vista uses MLDv2 by default, but will use MLD if it receives an MLD message. You can configure IPv6 to use MLD with the “**netsh interface ipv6 set global mldversion=version2**” command.

Querier Status

Shows the Querier status as “**ACTIVE**” or “**IDLE**”. Displays “DISABLE” to note that 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.

The MLD Snooping Status table **Router Port** parameters are explained below:

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

LIB-44xx port number (e.g., ports 1-4 on the LIB-4400).

Note: the FPGA port (port 12 or LIB-4424 port 24) is “hidden” when the Shared port is set to Internal mode.

Status

Indicate whether or not a specific port is a router port.

Buttons

Refresh: Click to refresh the page immediately.

Clear: Clears all Statistics counters.

Auto-refresh: Check this box to enable an automatic refresh of the page at 3 second intervals.

MLD Snooping > Groups Information

The **MLD Snooping > Groups Information** menu path displays entries in the MLD Group Table.

The MLD Groups Table is sorted first by VLAN ID, and then by group.

Each page shows up to 99 entries from the MLD Group table selected through the "**entries per page**" input field (the default is 20). When first visited, the web page will show the first 20 entries from the beginning of the MLD Group Table.

The "Start from VLAN", and "group address" input fields let you select the starting point in the MLD Group Table. Click the **Refresh** button to update the displayed table starting from that or the next closest MLD Group Table match. In addition, the two input fields will - upon a **Refresh** button click - assume the value of the first displayed entry, allowing for continuous refresh with the same start address.

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

The **MLD Snooping > Groups Information** parameters are explained below:

VLAN ID

The VLAN ID of the entry (e.g., VLAN ID 1).

Groups

Name of the Group displayed.

Port Members

Ports under this group (e.g., ports 1-4 on the LIB-4400).

Note: the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Buttons

Auto-refresh: Check this checkbox to enable an automatic refresh of the page at 3 second intervals.

Refresh: Refreshes the displayed table starting from the input fields.

|<<: Updates the table starting from the first entry in the MLD Group Table.

>>: Updates the table, starting with the entry after the last entry currently displayed.

MLD Snooping > IPv6 SFM Information

The **MLD Snooping > IPv6 SFM Information** menu path displays entries in the MLD SFM Information Table are shown on this page. The MLD SFM (Source-Filtered Multicast) Information Table also contains the SSM (Source-Specific Multicast) information. This table is sorted first by VLAN ID, then by group, and then by Port. Different source addresses belong to the same group are treated as single entry.

Each page shows up to 99 entries from the MLD SFM Information 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 MLD SFM Information Table.

The "Start from VLAN", and "Group" input fields allow the user to select the starting point in the MLD SFM Information Table. Clicking the **Refresh** button will update the displayed table starting from that or the closest next MLD SFM Information Table match. In addition, the two input fields will - upon a **Refresh** button click - assume the value of the first displayed entry, allowing for continuous refresh with the same start address.

The >> button will use the last entry of the currently displayed as a basis for the next lookup. When the end is reached, the text "No more entries" displays in the displayed table. Use << the button to start over.

VLAN ID

VLAN ID of the group.

Group

The Group address of the group displayed (e.g., **ff02::1:ff00:108**).

Port

The LIB-44xx port number (e.g., **1 - 4**).

Mode

Indicates the filtering mode maintained per (VLAN ID, port number, Group Address) basis. It can be either **'Include'** or **'Exclude'**.

Source Address

The IP Address of the source. The LIB-44xx limits the total number of IP source addresses for filtering to 128.

Type

Indicates the type of action. It can be either **'Allow'** or **'Deny'**.

Hardware Filter/Switch

Indicates whether the data plane destined to the specific group address from the source IPv6 address could be handled by chip.

Buttons

Auto-refresh: Check this box to enable an automatic refresh of the page at 3 second intervals.

Refresh: Refreshes the displayed table starting from the input fields.

|<<: Updates the table starting from the first entry in the MLD SSM Information Table.

>>: Updates the table, starting with the entry after the last entry currently displayed.

Monitor > LLDP

The **Monitor > LLDP** menu path displays the **Neighbours**, **LLDP Neighbours** and **Port Statistics** sub-menus.

The IEEE 802.1ab Link Layer Discovery Protocol (LLDP) 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).

Monitor > LLDP > Neighbours

The **Monitor > LLDP > Neighbours** menu path provides a status overview of all LLDP neighbours.

Local Interface	Chassis ID	Port ID	Port Description	System Name	System Capabilities	Management Address
GigabitEthernet 1/21	04-BF-6D-1C-95-28	13		Switch09	Bridge(+)	172.16.20.9 (IPv4)
GigabitEthernet 1/21	172.16.20.77	00-15-65-A9-D6-7D	WAN PORT	W52P	Bridge(-), Router(+), Telephone(+)	

The displayed table contains a row for each port on which an LLDP neighbour is detected. If no neighbours are detected, the message “No neighbour information found” displays.

The columns hold the following information:

Local Port

The port on which the LLDP frame was received (e.g., **Port 3** on the screen above).

Chassis ID

The **Chassis ID** is the identification of the neighbour's LLDP frames (e.g., **00-C0-F2-00-00-01**).

Port ID

The Remote **Port ID** is the identification of the neighbour port (e.g., **1001**).

Port Description

The port description being advertised by the neighbour unit.

System Name

System Name is the name advertised by the neighbour unit.

System Description

The description advertised by the neighbour unit.

System Capabilities

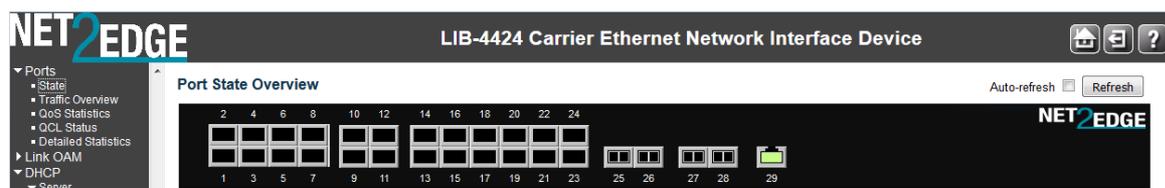
System Capabilities describes the neighbour unit's capabilities. The possible capabilities are:

1. **Other**: capabilities other than those listed below:
2. **Repeater**: the neighbour unit functions with repeater capabilities.
3. **Bridge**: the neighbour unit functions with bridge capabilities (e.g., **Bridge(+)** shown above).
4. **WLAN Access Point (WAP)**: the neighbour unit functions with WAP capabilities.
5. **Router**: the neighbour unit functions with router capabilities.
6. **Telephone**: the neighbour unit functions with telephone capabilities.
7. **DOCSIS cable device**: the neighbour unit functions with DOCSIS capabilities.
8. **Station only**: the neighbour unit functions with just station capabilities.
9. **Reserved**: this neighbour unit function description is reserved for future use.

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 neighbour unit's address that is used for higher layer entities to assist discovery by the network management. This could for example hold the neighbour's IP address. When you click the **Management Address** link (e.g., [192.168.1.10 \(IPv4\)](#) in the screen example above), a security dialog opens. Enter the valid password information to display the neighbour's startup screen.



The neighbour device is a Net2Edge Switch in the example above. You can click the browser Back button to return to the LIB-44xx LLDP Neighbour Information page.

At the switch, you can also view the Port Statistics from the **Monitor > LLDP** menu path, and click the **Management Address** link to display the INDURA neighbour's startup screen (i.e., go back to the LIB-44xx startup screen display).

Buttons

Refresh: Click to refresh the page immediately.

Auto-refresh: Check this box to enable an automatic refresh of the page at 3 second intervals.

Monitor > LLDP > Port Statistics

The **Monitor > LLDP > Port Statistics** menu path provides an overview of all LLDP traffic.

Two types of counters are shown. **Global Counters** are counters that refer to the LIB-44xx at the device **level**, while **Local Counters** refer to per port counters for the currently selected LIB-44xx.

The screenshot displays the NET2EDGE web interface for a LIB-4424 Carrier Ethernet Network Interface Device. The left sidebar shows a navigation menu with categories like Groups, Information, IPv4 SFM, IPv6 SFM, LLDP, Ethernet Services, and Performance. The main content area is titled 'LLDP Global Counters' and includes a 'Global Counters' table and an 'LLDP Statistics Local Counters' table.

LLDP Global Counters

Global Counters	
Clear global counters	<input checked="" type="checkbox"/>
Neighbor entries were last changed 2016-11-30T11:26:16+00:00 (360222 secs. ago)	
Total Neighbors Entries Added	4
Total Neighbors Entries Deleted	2
Total Neighbors Entries Dropped	0
Total Neighbors Entries Aged Out	2

LLDP Statistics Local Counters

Local Interface	Tx Frames	Rx Frames	Rx Errors	Frames Discarded	TLVs Discarded	TLVs Unrecognized	Org. Discarded	Age-Outs	Clear
*	*	*	*	*	*	*	*	*	<input checked="" type="checkbox"/>
GigabitEthernet 1/1	0	0	0	0	0	0	0	0	<input checked="" type="checkbox"/>
GigabitEthernet 1/2	0	0	0	0	0	0	0	0	<input checked="" type="checkbox"/>
GigabitEthernet 1/3	0	0	0	0	0	0	0	0	<input checked="" type="checkbox"/>
GigabitEthernet 1/4	0	0	0	0	0	0	0	0	<input checked="" type="checkbox"/>
GigabitEthernet 1/5	0	0	0	0	0	0	0	0	<input checked="" type="checkbox"/>
GigabitEthernet 1/6	0	0	0	0	0	0	0	0	<input checked="" type="checkbox"/>

The **Port Statistics** page parameters are explained below:

LLDP Global Counters

Neighbour entries were last changed on

Shows the date and/or time when the last entry was last deleted or added. It also shows the time elapsed since the last change was detected.

Total Neighbours Entries Added

Shows the number of new entries added since the last LIB-44xx reboot.

Total Neighbours Entries Deleted

Shows the number of new entries deleted since the last LIB-44xx reboot.

Total Neighbours Entries Dropped

Shows the number of LLDP frames dropped due to the entry table being full.

Total Neighbours Entries Aged Out

Shows the number of entries deleted due to Time-To-Live (TTL) expiring.

LLDP Statistics Local Counters

The displayed table contains a row for each port. The columns hold the following information:

Local Port

The port on which LLDP frames are received or transmitted (e.g., ports 1-4 on the LIB-4400).

Note: the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Tx Frames

The number of LLDP frames transmitted on the port.

Rx Frames

The number of LLDP frames received on the port.

Rx Errors

The number of received LLDP frames containing some kind of error.

Frames Discarded

If an LLDP frame is received on a port, and the LIB-44xx's internal table has run full, the LLDP frame is counted and discarded. This situation is known as "Too Many Neighbours" 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 port's link is down, an LLDP shutdown frame is received, or when the entry ages out. See IEEE Std 802.1AB for more information.

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

The number of organizationally received TLVs.

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.

Buttons

Refresh: Click to refresh the page immediately.

Clear: Clears the **local counters**. If you click the **OK** button at the webpage confirmation dialog, all counters (including **global counters**) are cleared upon reboot.

Auto-refresh: Check this checkbox to enable an automatic refresh of the page at 3 second intervals.

Monitor > Ethernet Services> EVC Statistics

The **Monitor > Ethernet Services > EVC Statistics** menu path provides NNI port traffic statistics for the selected EVC. It also shows counters for UNI ports of ECEs mapping to the EVC.

Service Frame (Traffic) Colors - Green / Yellow / Red

The MEF specifies traffic “coloring” as a way to mark traffic as ‘in profile’ or ‘out of profile’ as it leaves the ingress UNI. MEF 10 specifies three levels of Bandwidth Profile compliance:

Green: Service Frame subject to SLA; in-profile and conform to BW profile; delivered per the service performance objectives specified. A service frame is green if it is conformant with the CIR of the bandwidth profile.

Yellow: Service Frame not subject to SLA; out of profile but typically not immediately discarded; not delivered per the service performance objectives; may get discarded by the network. A service frame is yellow if it is not conformant with the EIR of the bandwidth profile.

Red: Service Frame discarded; out of profile and immediately discarded. A service frame is red if it is conformant with neither the CIR nor EIR of the bandwidth profile.

The EVC Statistics table can be set to display statistics in terms of frames, or bytes, or both. The default EVC Statistics table displays with “Frames” as the unit of measure, as shown below:

Clear	EVC ID	CoS ID	Green Frames		Yellow Frames		Red Frames		Discarded Frames	
			RX	TX	RX	TX	RX	TX	RX	TX
<input type="checkbox"/>	1	0	0	0	0	0	0	0	0	0

The EVC Statistics table parameters are explained below:

Clear

This box marks a port for clearance in next Clear operation.

Port

The UNI/NNI port for the EVC (e.g., ports 1-4 on the LIB-4400).

Note: the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Rx Green

The number of green frames received.

Tx Green

The number of green frames transmitted.

Rx Yellow

The number of yellow frames received.

Tx Yellow

The number of yellow frames transmitted.

Rx Red

The number of red received.

Rx Discarded

The number of discarded in the ingress queue system.

Tx Discarded

The number of discarded in the egress queue system.

Buttons

Frames: Radio button that shows frames statistics only.

Bytes: Radio button that shows bytes statistics only.

Both: Radio button that shows both frames and bytes statistics.



: Use the port select box to select which port is affected by clicking the buttons.

Auto-refresh: Check this checkbox to enable an automatic refresh of the page at 3 second intervals.

Refresh: Click to refresh the page immediately.

Clear: Clears the counters for selected ports.

Clear: Clears the EVC statistics counters for selected ports. Displays the message “All EVCs statistics will be cleared. Do you want to proceed anyway?”.



Click the **OK** button to clear (zero out) the EVC statistics counters for selected ports, or click the **Cancel** button to clear the webpage message but leave the EVC statistics counters the same.

Example

The EVC Statistics table below is set to display statistics with both frames and bytes as the unit of measure.

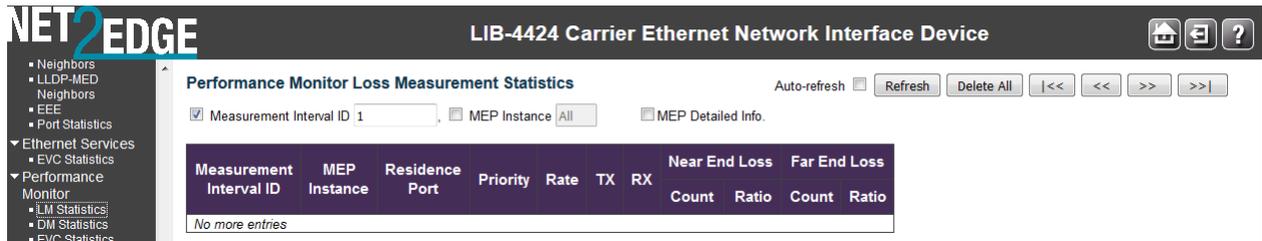
Clear	EVC ID	CoS ID	Green		Yellow		Red		Discarded					
			Frames	Bytes	Frames	Bytes	Frames	Bytes	Frames	Bytes				
			RX	TX	RX	TX	RX	TX	RX	TX				
<input type="checkbox"/>	1	0	0	0	0	0	0	0	0	0	0	0	0	0

Monitor > Performance Monitor (PM)

Performance Monitor will continuously check the bandwidth usage of each instance, if a threshold crossing alarm condition occurs, the system will generate TCA traps to server while SNMP trap is enabled. Up to 132 Instances are supported.

> LM Statistics

This page provides the performance monitor loss measurement traffic statistics for the selected measurement interval ID and Loss Measurement instance. The screens below show the statistics page with and without the **MEP Detailed Info** checkbox checked.



The parameters are described below:

Measurement Interval ID

The measurement interval for the performance monitor data sets.

MEP Instance

The MEP instance for the performance monitor data sets.

Residence Port

The residence port for the MEP.

Priority

The priority to be inserted as PCP bits in TAG (if any).

Rate

The gap between transmitting 1DM/DMM PDUs in 10ms increments. The range is 10 - 65535.

Unit

The time resolution.

TX

The number of transmitted.

RX

The number of received.

Near End Loss Count

The near end loss count.

Near End Loss Ratio

The near end loss ratio.

Far End Loss Count

The far end loss count.

Far End Loss Ratio

The far end loss ratio.

Domain

Port: This is a MEP in the Port Domain. 'Flow Instance' is a Port.

Evc: This is a MEP in the EVC Domain. 'Flow Instance' is a EVC.

VLAN: This is a MEP in the VLAN Domain. 'Flow Instance' is a VLAN.

Direction

Up: This is a Down MEP - monitoring ingress OAM and traffic on 'Residence Port'.

Down: This is a Up MEP - monitoring egress OAM and traffic on 'Residence Port'.

Level

The MEG level of this MEP.

Flow Instance

The MEP is related to this flow - See 'Domain'.

Tagged VID

Port MEP: An outer C/S-tag (depending on VLAN Port Type) is added with this VID. Entering '0' means no TAG added.

EVC MIP: On Serval, this is the Subscriber VID that identify the subscriber flow in this EVC where the MIP is active.

MEP ID

This value will become the transmitted two byte CCM MEP ID.

MAC Address

The MAC of this MEP - can be used by other MEP when unicast is selected (Info only).

Peer MEP ID

This value will become an expected MEP ID in a received CCM - see 'cMEP'.

Peer MAC Address

This MAC will be used when unicast is selected with this peer MEP. Also this MAC is used to create HW checking of receiving CCM PDU (LOC detection) from this MEP.

Buttons

One-way: Show one-way statistics only.



Two-way: Show two-way statistics only.



Both: Show both frames and bytes statistics.

Auto-refresh: Check this box to refresh the page automatically every three seconds.

Refresh: Click to refresh the page immediately.

Delete All: Delete all table entries.

|<<: Updates the table entries, starting from the first available entry.

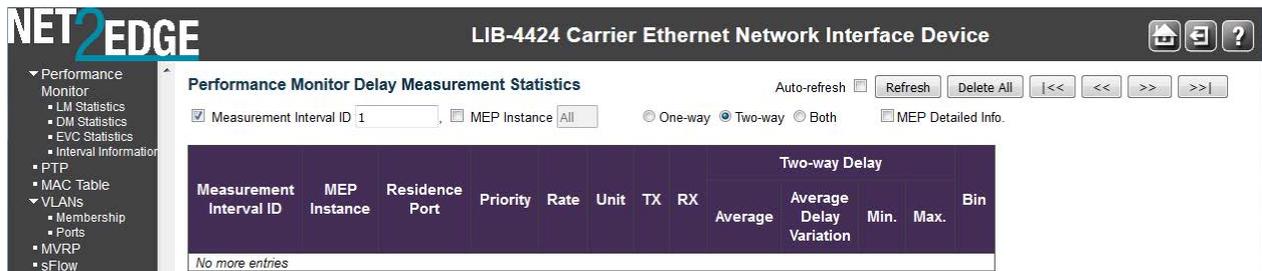
<<: Updates the table entries, ending at the last entry currently displayed.

>>: Updates the table entries, starting from the last entry currently displayed.

>>|: Updates the table entries, ending at the last available entry.

> DM Statistics

This page provides the performance monitor delay measurement traffic statistics for the selected measurement interval ID and Delay Measurement instance. The menu path is **Monitor > Performance Monitor > DM Statistics**.



The delay measurement traffic statistics parameters are described below:

Measurement Interval ID

The measurement interval for the performance monitor data sets.

MEP Instance

The MEP instance for the performance monitor data sets.

Residence Port

The residence port for the MEP.

Priority

The priority to be inserted as PCP bits in TAG (if any).

Rate

Selected the frame rate of CCM/LMM PDU. This is the inverse of transmission period as described in Y.1731.

Unit

The time resolution.

TX

The number of transmitted.

RX

The number of received.

One-way Far to Near Average Delay

The one-way far to near average delay.

One-way Far to Near Average Delay Variation

The one-way far to near average delay variation.

One-way Far to Near Min. Delay

The minimum one-way near to far delay.

One-way Far to Near Max. Delay

The maximum one-way near to far delay.

One-way Near to Far Average Delay

The number of red received.

One-way Near to Far Average Delay Variation

The one-way near to far average delay variation.

One-way Near to Far Min. Delay

The minimum one-way near to far delay.

One-way Near to Far Max. Delay

The maximum one-way near to far delay.

Two-way Delay Average Delay

The two-way average delay.

Two-way Average Delay Variation

The two-way average delay variation.

Two-way Min. Delay

The minimum two-way delay.

Two-way Max. Delay

The maximum two-way delay.

Domain

Port: This is a MEP in the Port Domain. 'Flow Instance' is a Port.

Evc: This is a MEP in the EVC Domain. 'Flow Instance' is a EVC.

VLAN: This is a MEP in the VLAN Domain. 'Flow Instance' is a VLAN.

Direction

Up: This is a Down MEP - monitoring ingress OAM and traffic on 'Residence Port'.

Down: This is a Up MEP - monitoring egress OAM and traffic on 'Residence Port'.

Level

The MEG level of this MEP.

Flow Instance

The MEP is related to this flow - See 'Domain'.

Tagged VID

Port MEP: An outer C/S-tag (depending on VLAN Port Type) is added with this VID. Entering '0' means no TAG added.

EVC MIP: On Serval, this is the Subscriber VID that identify the subscriber flow in this EVC where the MIP is active.

MEP ID

This value will become the transmitted two byte CCM MEP ID.

MAC Address

The MAC of this MEP - can be used by other MEP when unicast is selected (Info only).

Peer MEP ID

This value will become an expected MEP ID in a received CCM - see 'cMEP'.

Peer MAC Address

This MAC will be used when unicast is selected with this peer MEP. Also this MAC is used to create HW checking of receiving CCM PDU ([LOC](#) detection) from this MEP.

Buttons

Auto-refresh: Check this box to refresh the page automatically every **three seconds**.

Refresh: Click to refresh the page immediately.

Delete All: Delete all table entries.

|<<: Updates the table entries, starting from the first available entry.

<<: Updates the table entries, ending at the last entry currently displayed.

>>: Updates the table entries, starting from the last entry currently displayed.

>>|: Updates the table entries, ending at the last available entry.

> EVC Statistics

This page provides the performance monitor EVC traffic statistics for the selected measurement interval ID and EVC instance from the **Monitor > Performance Monitor > EVC Statistics** menu path. In order to get EVC statistics into the PM data set, an EVC must be created. The PM data set contains EVC statistics for all created EVC(s) for each configured port (up to the max data set limit). The EVC Statistic PM data set must contain per UNI and NNI Port related to this EVC instance information. The "Total" count on a port is the sum of all the ECE counts related to this EVC on that port.

The screenshot shows the 'Performance Monitor EVC Statistics' page. At the top, there are controls for 'Auto-refresh' (unchecked), 'Refresh', 'Delete All', and navigation arrows. Below these are input fields for 'Measurement Interval ID' (set to 1) and 'EVC Instance' (set to All). There are radio buttons for 'Frames' (selected), 'Bytes', and 'Both'. The table below has the following structure:

Measurement Interval ID	EVC Instance	Port	Cos	Green Frames		Yellow Frames		Red Frames		Discarded Frames	
				Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx
No more entries											

The Performance Monitor EVC traffic Statistics parameters are described below:

Measurement Interval ID

The measurement interval for the performance monitor data sets.

EVC Instance

The EVC instance for the performance monitor data sets.

MEP Instance

The MEP instance for the performance monitor data sets.

Residence Port

The residence port for the EVC.

Rx Green Frames

The number of green received.

Tx Green Frames

The number of green transmitted.

Rx Yellow Frames

The number of yellow received.

Tx Yellow Frames

The number of yellow transmitted.

Rx Red Frames

The number of red received.

Rx Discarded Frames

The number of discarded in the ingress queue system.

Tx Discarded Frames

The number of discarded in the egress queue system.

Buttons

Frames: Show frames statistics only.

Bytes: Show bytes statistics only.

Both: Show both frames and bytes statistics.

Auto-refresh: Check this box to refresh the page automatically every three seconds.

Refresh: Click to refresh the page immediately.

Delete All: Delete all table entries.

|<<: Updates the table entries, starting from the first available entry.

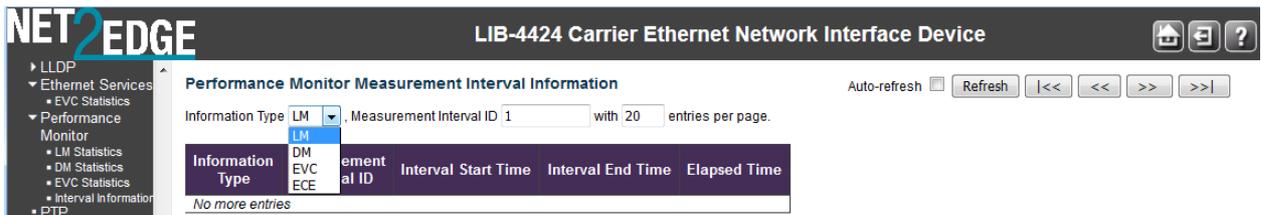
<<: Updates the table entries, ending at the last entry currently displayed.

>>: Updates the table entries, starting from the last entry currently displayed.

>>|: Updates the table entries, ending at the last available entry.

> Interval Information

This page provides the performance monitor measurement interval information from the **Monitor > Performance Monitor > Interval Information** menu path.



This page provides the performance monitor measurement interval information.

Information Type

The type for the performance monitor data sets. The valid types are:

LM: Loss Measurement type selected.

DM: Delay Measurement type selected.

EVC: This PM instance is per-port per-EVC based. The specific port is configured by the Port field; the EVC is configured by the EVC field.

ECE: This PM instance is per port per ECE based. The specific port is configured by the Port field; the ECE is configured by the ECE field.

Port: This PM instance is port based. The specific port is configured by the Port field.

Measurement Interval ID

The measurement interval for the performance monitor data sets.

Interval Start Time

The PM interval starting day, date, and time.

Interval End Time

The interval ending day, date, and time.

Elapsed Time

The PM elapsed time in seconds.

Buttons

Auto-refresh: Check this box to refresh the page automatically. Automatic refresh occurs every **three seconds**.

Refresh: Click to refresh the page immediately.

<<: Updates the table entries, starting from the first available entry.

<<: Updates the table entries, ending at the last entry currently displayed.

>>: Updates the table entries, starting from the last entry currently displayed.

>>|: Updates the table entries, ending at the last available entry.

Examples

Port: Performance monitor measurement interval data for the **Port** information type:

Performance Monitor Measurement Interval Information Auto-refresh Refresh <<< << >> >>|

Information Type **Port**, Measurement Interval ID with entries per page.

Information Type	Measurement Interval ID	Interval Start Time	Interval End Time	Elapsed Time
Port	190	1970-01-06T21:40:57+00:00	1970-01-06T21:55:57+00:00	900
Port	191	1970-01-06T21:55:57+00:00	1970-01-06T22:10:57+00:00	900
Port	192	1970-01-06T22:10:57+00:00	1970-01-06T22:25:57+00:00	900
Port	193	1970-01-06T22:25:58+00:00	1970-01-06T22:40:58+00:00	900
Port	194	1970-01-06T22:40:58+00:00	1970-01-06T22:55:58+00:00	900
Port	195	1970-01-06T22:55:58+00:00	1970-01-06T23:10:58+00:00	900
Port	196	1970-01-06T23:10:58+00:00	1970-01-06T23:25:58+00:00	900
Port	197	1970-01-06T23:25:58+00:00	1970-01-06T23:40:58+00:00	900
Port	198	1970-01-06T23:40:59+00:00	1970-01-06T23:55:59+00:00	900
Port	199	1970-01-06T23:55:59+00:00	1970-01-07T00:10:59+00:00	900
Port	200	1970-01-07T00:10:59+00:00	1970-01-07T00:25:59+00:00	900
Port	201	1970-01-07T00:25:59+00:00	1970-01-07T00:40:59+00:00	900
Port	202	1970-01-07T00:40:59+00:00	1970-01-07T00:55:59+00:00	900
Port	203	1970-01-07T00:55:59+00:00	1970-01-07T01:10:59+00:00	900
Port	204	1970-01-07T01:11:00+00:00	1970-01-07T01:26:00+00:00	900
Port	205	1970-01-07T01:26:00+00:00	1970-01-07T01:41:00+00:00	900
Port	206	1970-01-07T01:41:00+00:00	1970-01-07T01:56:00+00:00	900
Port	207	1970-01-07T01:56:00+00:00	1970-01-07T02:11:00+00:00	900
Port	208	1970-01-07T02:11:00+00:00	1970-01-07T02:26:00+00:00	900
Port	209	1970-01-07T02:26:01+00:00	1970-01-07T02:41:01+00:00	900

EVC: Performance monitor measurement interval data for the **EVC** information type:

Performance Monitor Measurement Interval Information Auto-refresh Refresh <<< << >> >>|

Information Type **EVC**, Measurement Interval ID with entries per page.

Information Type	Measurement Interval ID	Interval Start Time	Interval End Time	Elapsed Time
EVC	190	1970-01-06T21:40:57+00:00	1970-01-06T21:55:57+00:00	900
EVC	191	1970-01-06T21:55:57+00:00	1970-01-06T22:10:57+00:00	900
EVC	192	1970-01-06T22:10:57+00:00	1970-01-06T22:25:57+00:00	900
EVC	193	1970-01-06T22:25:58+00:00	1970-01-06T22:40:58+00:00	900
EVC	194	1970-01-06T22:40:58+00:00	1970-01-06T22:55:58+00:00	900
EVC	195	1970-01-06T22:55:58+00:00	1970-01-06T23:10:58+00:00	900
EVC	196	1970-01-06T23:10:58+00:00	1970-01-06T23:25:58+00:00	900
EVC	197	1970-01-06T23:25:58+00:00	1970-01-06T23:40:58+00:00	900
EVC	198	1970-01-06T23:40:58+00:00	1970-01-06T23:55:58+00:00	900
EVC	199	1970-01-06T23:55:59+00:00	1970-01-07T00:10:59+00:00	900
EVC	200	1970-01-07T00:10:59+00:00	1970-01-07T00:25:59+00:00	900
EVC	201	1970-01-07T00:25:59+00:00	1970-01-07T00:40:59+00:00	900
EVC	202	1970-01-07T00:40:59+00:00	1970-01-07T00:55:59+00:00	900
EVC	203	1970-01-07T00:55:59+00:00	1970-01-07T01:10:59+00:00	900
EVC	204	1970-01-07T01:11:00+00:00	1970-01-07T01:26:00+00:00	900
EVC	205	1970-01-07T01:26:00+00:00	1970-01-07T01:41:00+00:00	900
EVC	206	1970-01-07T01:41:00+00:00	1970-01-07T01:56:00+00:00	900
EVC	207	1970-01-07T01:56:00+00:00	1970-01-07T02:11:00+00:00	900
EVC	208	1970-01-07T02:11:00+00:00	1970-01-07T02:26:00+00:00	900
EVC	209	1970-01-07T02:26:00+00:00	1970-01-07T02:41:00+00:00	900

Note: the Performance monitor measurement interval data for the **ECE** information type is very similar.

Sample Stored Report

A portion of the unzipped “_history_1970-01-01_00h.15m.30s_DM_all.csv.gz” file is shown below (in .Xls file format).

For output the range is 1-25000000 Hz (1 Hz to 25 MHz).

Actual Frequency

Displays the real time frequency detected into the SMB input. Active even if the SMB input is disabled (not being used internally by the device).

External I/O Options

Impedance

Shows the current Impedance used by the External Clock. The valid values are:

50 Hz : Enable 50 ohm impedance.

75 Hz : Enable 75 ohm impedance.

Hi-Z : No impedance termination driven, "tri-stated" or "floating".

undefined: No impedance configuration has been set.

PTP Clock Description

Clock Instance

Indicates the Instance of a particular Clock Instance (**0-3**). Click on the Clock Instance number to monitor that Clock's details.

Device Type

Indicates the Type of the Clock Instance. There are five 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.

MastrOnly: Clock's Device Type is Master Only.

slaveOnly: Clock's Device Type is Slave Only.

Port List

Shows the ports configured for the Clock instance with a green check mark () (e.g., ports 1-4 on the LIB-4400).

Note: the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Buttons

Auto-refresh: Check this checkbox to enable automatic page refreshes at 3 second intervals.

Refresh: Click to refresh the page immediately.

PTP Clock's Configuration

Click on a Clock Instance in the PTP Clock's Configuration section at the **Monitor > PTP** menu path to display that **PTP Clock's Configuration** page.

PTP Clock's Configuration

Auto-refresh Refresh

Clock Instance	HW Domain	Device Type	Profile	Filter Type
0	0	Ord-Bound	1588	ACI_BC_FULL_ON_PATH_FREQ

Local Clock Current Time

PTP Time	Clock Adjustment method	Ports Monitor Page
1970-01-10T02:34:09+00:00 000,776,900	Internal Timer	Ports Monitor

Clock Default DataSet

Device Type	One-Way	2 Step Flag	Ports	Clock Identity	Dom	Clock Quality	Pri1	Pri2	Local Prio	Protocol	VID	PCP	DSCP
Ord-Bound	False	False	29	00:01:c1:ffe:00:fc:b0	0	Cl:248 Ac:Unknwn Va:65535	128	128	128	Ethernet	1	0	0

Clock Current DataSet

stpRm	Offset From Master	Mean Path Delay	Slave Port	Slave State	Holdover(ppb)
0	0.000,000,000	0.000,000,000	0	FREERUN	N.A.

Clock Parent DataSet

Parent Port Identity	Port	PStat	Var	ChangeRate	Grand Master Identity	Grand Master Clock Quality	Pri1	Pri2
00:01:c1:ffe:00:fc:b0	0	False	0	0	00:01:c1:ffe:00:fc:b0	Cl:248 Ac:Unknwn Va:65535	128	128

Clock Time Properties DataSet

UtcOffset	Valid	leap59	leap61	Time Trac	Freq Trac	ptp Time Scale	Time Source
0	False	False	False	False	False	True	160

The **Local Clock Current Time** table displays the current PTP Time (e.g., 1970-01-03T19:52:44+00:00 541,523,780), Clock Adjustment method (e.g., Internal Timer) and [Ports Monitor Page](#) link. When you click the [Ports Monitor](#) link, the PTP Clock's Port Data Set Configuration page displays (see below).

The **Clock Default DataSet** table displays the current ClockId (e.g., 0) Device Type (e.g., P2pTransp) 2 Step Flag (e.g., True) Ports (e.g., 9) Clock Identity (e.g., 00:c0:f2:ff:fe:00:00:01) Dom (e.g., 0) Clock Quality (e.g., Cl:251 Ac:254 Va:65535) Pri1 (e.g., 128) Pri2 (e.g., 128) Protocol (e.g., Ethernet) One-Way (e.g., False) VLAN Tag Enable (e.g., False) VID (e.g., 0) and PCP (e.g., 0) information.

The clock default data set is defined in the IEEE 1588 Standard. It holds three groups of data: the static members defined at clock creation time, the dynamic members defined by the system, and configurable members which can be set here.

The **Clock Current DataSet** table displays the current stpRm (e.g., 0) Offset From Master (e.g., 0.000,000,000) and Mean Path Delay (e.g., 0.000,000,000) information. The clock current data set is defined in the IEEE 1588 Standard. The current data set is dynamic.

The **Clock Parent DataSet** table displays the current Parent Port Identity (e.g., 00:c0:f2:ff:fe:00:00:01) Port (e.g., 0) PStat (e.g., False) Var (e.g., 0) ChangeRate (e.g., 0) Grand Master Identity (e.g., 00:c0:f2:ff:fe:00:00:01) Grand Master Clock Quality (e.g., Cl:251 Ac:254 Va:65535) Pri1 (e.g., 128) and Pri2 (e.g., 128) information. The clock current data set is defined in the IEEE 1588 Standard. The current data set is dynamic.

The **Clock Time Properties DataSet** table displays the current UtcOffset (e.g., 0) Valid (e.g., False) leap59 (e.g., False) leap61 (e.g., False) Time Trac (e.g., False) Freq Trac (e.g., False) ptp Time Scale (e.g., True) and Time Source (e.g., 160) information.

The **Servo Parameters** table displays the current Display (e.g., False) P-enable (e.g., True) I-enable (e.g., True) D-enable (e.g., False) 'P' constant (e.g., 10) 'I' constant (e.g., 100) and 'D' constant (e.g., 1000) information.

The **Filter Parameters** table displays the current DelayFilter (e.g., 6) period (e.g., 1) and dist (e.g., 2) information.

The **Unicast Slave Configuration** table displays the current Index (e.g., 1) Duration (e.g., 100) IP_Address (e.g., 100.100.10.1) Grant (e.g., 0) and CommState (e.g., IDLE) information.

PTP Clock's Port Data Set Configuration

Click on the [Ports Monitor](#) Page link in the PTP Clock's Configuration section to display the **PTP Clock's Port Data Set Configuration** table. The port data set is defined in the IEEE 1588 Standard.

Port	Stat	MDR	PeerMeanPathDel	Anv	ATo	Syv	Dim	MPR	Delay Asymmetry	Ing. Latency	Egr. Latency	Version	Mcast Addr	Not Slave	Local Prio
1	dsbl	3	0.000,000,000	1	3	0	e2e	0	0.000,000,000	0.000,000,000	0.000,000,000	2	Default	False	128
2	dsbl	3	0.000,000,000	1	3	0	e2e	0	0.000,000,000	0.000,000,000	0.000,000,000	2	Default	False	128

The related **The PTP Clock's Port Data Set Configuration** table parameters are explained below:

Port: The LIB-44xx port number (e.g., 1-4). The Port number [1..max port no].

Stat: The Clock's port status (e.g., dsbl or p2pt). The current state of the port.

MDR: log Min Delay Req Interval: The delay request interval announced by the master (e.g., 0 or 3).

PeerMeanPathDel: The path delay measured by the port in P2P mode. In E2E mode this value is 0 (e.g., 0.000,000,000).

Anv: The interval for issuing Announce messages in master state (e.g., 1).

ATo: The timeout for receiving Announce messages on the port (e.g., 3).

Syv: The interval for issuing Sync messages in master (e.g., 0).

Dim: the port's delay mechanism (e.g., *dsbl* for disabled, or *p2pt* for Peer to Peer Transparent):

e2e: End to end delay measurement.

p2p: Peer to peer delay measurement.

dsbl: Delay is disabled.

MPR: The interval for issuing Delay_Req messages for the port in E2e mode. This value is announced from the master to the slave in an announce message. The value is reflected in the MDR field in the Slave. The interval for issuing Pdelay_Req messages for the port in P2P mode. **Note:** This value is interpreted as an interval (i.e., MPR = 0 => 1 Delay_Req pr sec) independent of the Sync rate (e.g., 3).

Delay Asymmetry: The transmission delay asymmetry for a link. See IEEE 1588 Section 7.4.2 Communication path asymmetry (e.g., 0.000,000,000).

Ingress Latency: Ingress latency measured in ns, as defined in IEEE 1588 Section 7.3.4.2. (e.g., 0.000,000,000).

Egress Latency: Egress latency measured in ns, as defined in IEEE 1588 Section 7.3.4.2. (e.g., 0.000,000,000).

Version: The current implementation only supports PTP version 2 (e.g., version 2).

Port Members

A green check mark (✓) indicates if a port is a member of the entry (CPU and Ports 1-4 and MGMT on the LIB-4400).

Note: the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Buttons

Auto-refresh: Check this checkbox to enable an automatic refresh of the page at 3 second intervals.

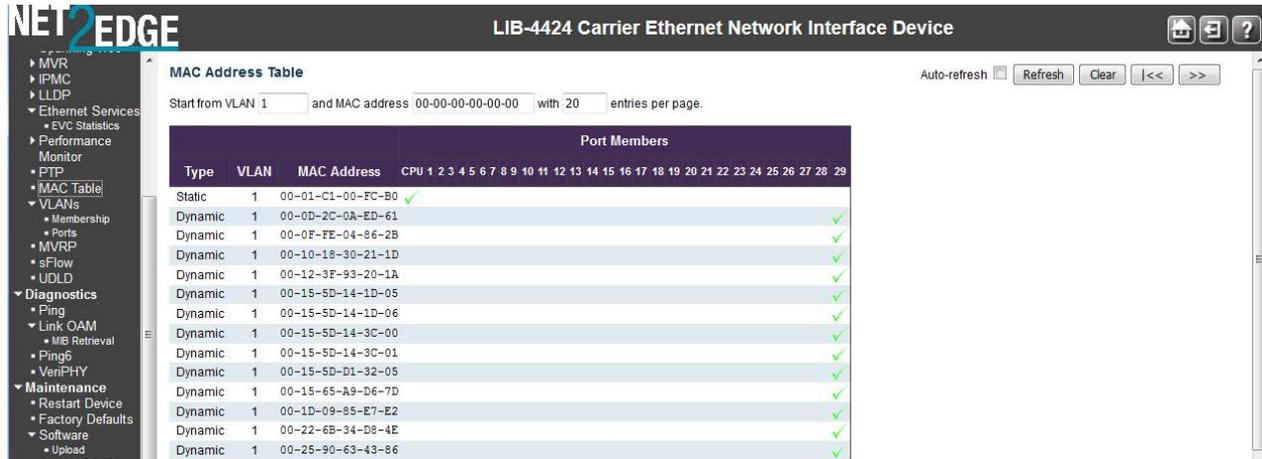
Refresh: Refreshes the displayed table starting from the "Start from MAC address" and "VLAN" input fields.

Clear: Flushes all dynamic entries.

<<: Updates the table starting from the first entry in the MAC Table (i.e., the entry with the lowest VLAN ID and MAC address).

>>: Updates the table, starting with the entry after the last entry currently displayed.

Example



Monitor > VLANs

The **Monitor > VLANs** menu path displays the **VLAN Membership** and **VLAN Port** sub-menus.

The LIB-44xx will be compliant with IEEE 802.1Q standard. The LIB-44xx is capable of VLAN bridging and filtering. By default the devices comes up with all ports belonging to the same VLAN.

The LIB-44xx supports the entire range of 4k VLAN IDs except for the following:

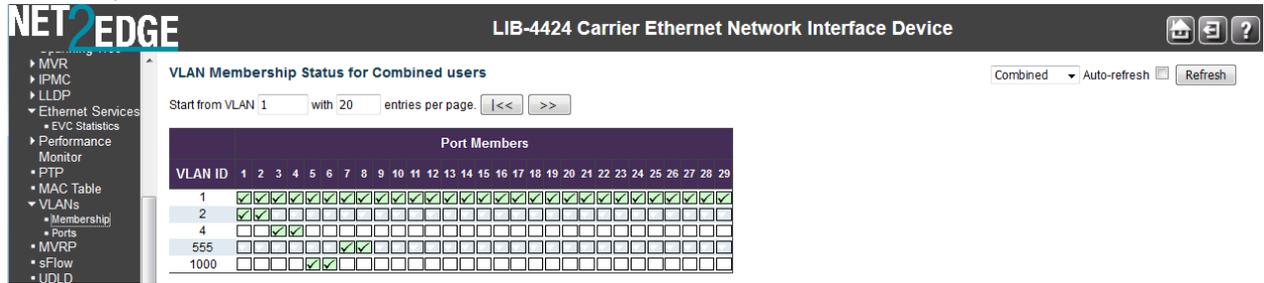
- a. VLAN ID = 0 is used for priority tagged traffic and will not be used.
- b. VLAN ID = 1 is used for the LIB-44xx default VLAN (native VLAN ID).
- c. VLAN ID = 4094 is reserved and not available for normal traffic.

Each VLAN has a unique string for identification called the "VLAN Name", no spaces will be allowed.

A Maximum of 64 VLANs can have VLAN names and the name is restricted to 32 bytes. Only alphabets and digits are allowed as valid characters with at least one alphabet to be part of the name.

Monitor > VLANs > VLAN Membership

The **Monitor > VLANs > VLAN Membership** menu path provides an overview of membership status of VLAN users. The default VLAN Membership Status page is shown below (for the 'Combined' selection).



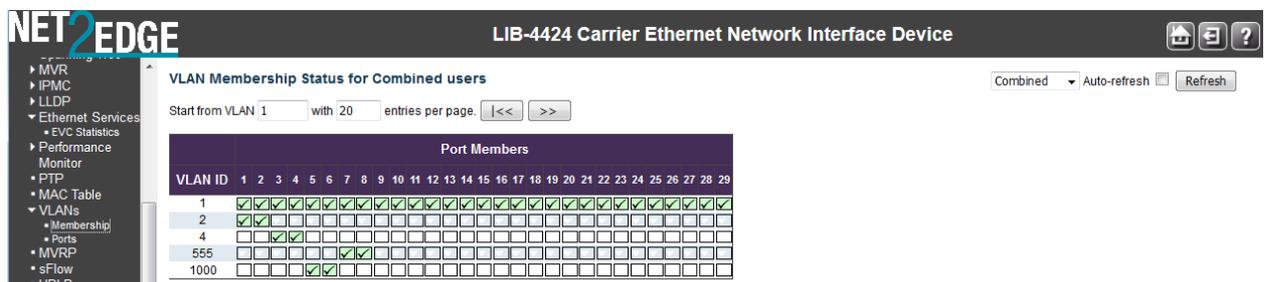
The VLAN Membership Status Page shows the current VLAN port members for all VLANs configured by a selected VLAN User (selection allowed by a combo box). When ALL VLAN Users are selected, it shows this information for all the VLAN Users, and this is by default. VLAN membership allows the frames classified to the VLAN ID to be forwarded on the respective VLAN member ports.

Each page displays up to 99 entries from the VLAN table (the default is 20 entries) selected via 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.

The "VLAN" input fields allow the user to select the starting point in the VLAN Table. Clicking the Refresh button will update the displayed table starting from that or the closest next VLAN Table match. The >> button uses the last entry of the currently displayed VLAN entry as a basis for the next lookup. When the end is reached the text "No more entries" displays in the displayed table. Use the << button to start over.

Clicking a Port Members checkbox alternately displays "VLAN included", then "VLAN not included" and then "Forbidden port" in the cursor over help (CoH).

The VLAN Membership Status page shown below has five VLAN IDs configured with various port members status.

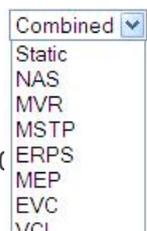


The **VLAN Membership** parameters are explained below:

VLAN USER

The VLAN User module uses the services of the VLAN management function to configure VLAN memberships and VLAN port configurations such as PVID and UVID. The LIB-44xx lets you monitor the following VLAN user types:

static: CLI/Web/SNMP are referred to as 'Static' VLAN user types.



NAS: NAS provides port-based authentication, which involves communications between a Supplicant, Authenticator, and an Authentication Server.

MVR: MVR is used to eliminate the need to duplicate multicast traffic for subscribers in each VLAN. Multicast traffic for all channels is sent only on a single (multicast) VLAN.

MSTP: The 802.1s Multiple Spanning Tree protocol (MSTP) uses VLANs to create multiple spanning trees in a network, which significantly improves network resource utilization while maintaining a loop-free environment.

ERPS: displays the status for all RPS (Ethernet Ring Protection Switching) user types, per ITU-T G.8032 Recommendation.

MEP: displays the status for all SOAM Maintenance End Point) user types.

EVC: displays the status for all EVC (Ethernet Virtual Circuit) user types.

VCL: displays the status for all VCL (VLAN Control List) user types.

Combined: displays the status for all of the above user types (the default).

Port Members

A row of check boxes for each port (e.g., ports 1-4 on the LIB-4400) is displayed for each VLAN ID.

✓ If a port is included in a VLAN, a check mark (✓) displays.

✗ If a port is included in a Forbidden port list, the image ✗ displays.

✗ If a port is included in a Forbidden port list and dynamic VLAN user register VLAN on the same Forbidden port, then the port in conflict displays with a ✗ image.

Note: the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

Buttons

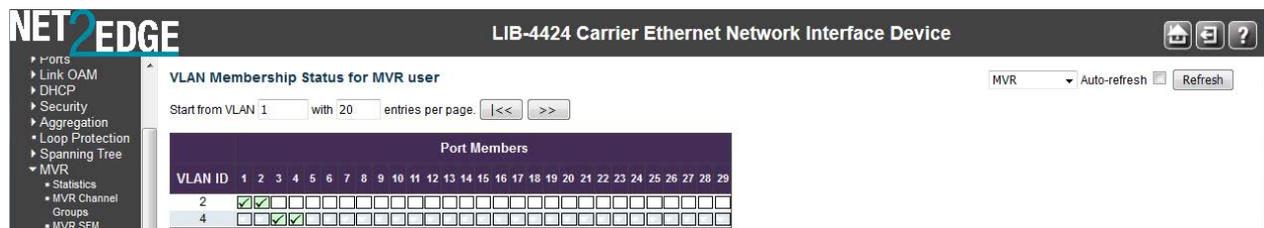
VLAN users dropdown : Select VLAN Users from this drop down list.

Auto-refresh: Check this box to enable an automatic refresh of the page every 3 seconds.

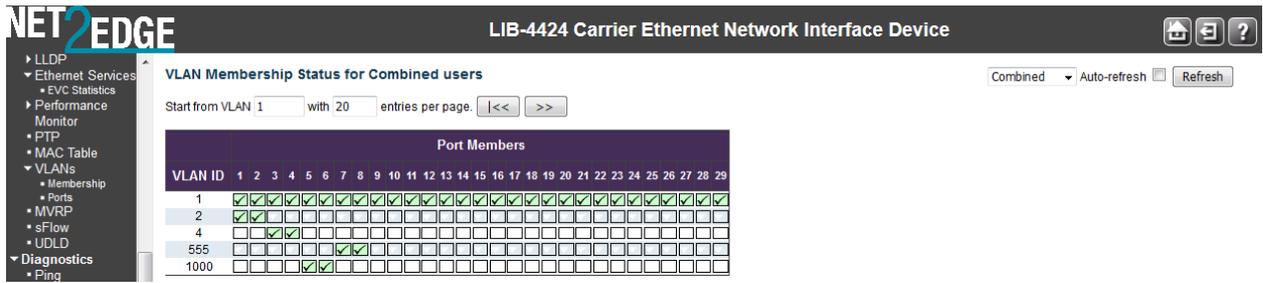
Refresh: Click to refresh the page immediately.

Examples

VLAN Membership Status for a **MVR** user (configured at the **Configuration > MVR** menu path):



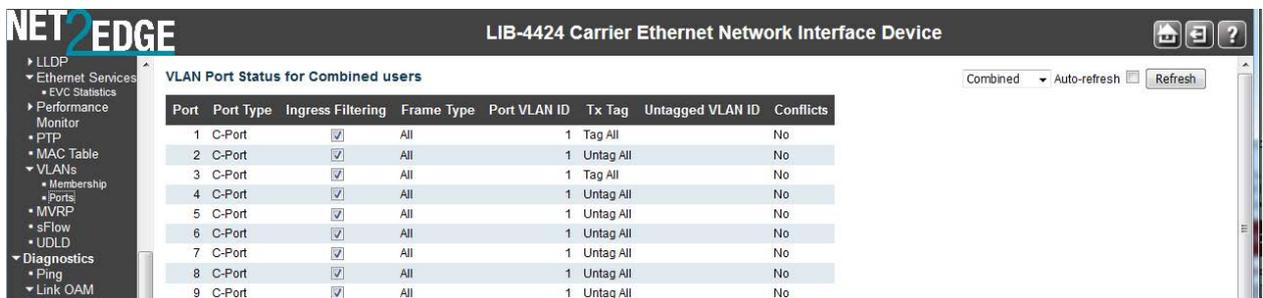
VLAN Membership Status for **Combined** (ALL) users:



Monitor > VLANs > VLAN Port

The **Monitor > VLANs > VLAN Port** menu path provides VLAN Port Status.

The LIB-44xx ports can be configured with a default or native VLAN id, so that all untagged and priority tagged traffic will be classified to this VLAN ID. The native VLAN ID for all ports is set to 1 but is configurable on a per-port basis. Hence all ports by default belong to the same broadcast domain.



The VLAN User module uses the services of the VLAN management functionality to configure VLAN memberships and VLAN port configuration such as PVID or UVID.

The **user type** combo box lets you select VLAN Users from this dropdown list, where:

Static: Displays only the static VLAN users' membership status. With a Static VLAN, assignments are created by assigning ports to a VLAN. As a device enters the network, the device automatically assumes the VLAN of the port. If you change ports and need access to the same VLAN, the network administrator must manually make a port-to-VLAN assignment for the new connection. CLI, Web, and SNMP are considered static.

NAS: Displays only the NAS VLAN users' membership status.

MVR: Displays only the MVR VLAN users' membership status.

MSTP: Displays only the MSTP VLAN users' membership status.

ERPS: Displays only the ERPS VLAN users' membership status.

MEP: Displays only the MEP VLAN users' membership status.

EVC: Displays only the EVC VLAN users' membership status.

VCL: Displays only the VCL VLAN users' membership status.

Combined: Displays the VLAN Membership Status for all of the possible user types (all of the above user types).



The VLAN Port Status parameters are explained below:

Port

The logical port (e.g., ports 1-4 for the LIB-4400) for the settings contained in the same row.

Note: the FPGA port (port 12 or LIB-4424 port 24) is "hidden" when the Shared port is set to Internal mode.

PVID

Displays the VLAN identifier for that port. The allowed values are **1** through **4094**. The default value is **1**.

Port Type

Displays the Port Type. Port type can be either Unaware, C-Port, S-Port, or Custom S-Port, where:

UnAware: all frames are classified to the Port VLAN ID and tags are not removed.

C-Port: a Customer Port.

S-Port: a Service port.

Custom S-Port: an S-port with Custom TPID.

Ingress Filtering

Shows the ingress filtering on a port. This parameter affects VLAN ingress processing. If ingress filtering is **Enabled** and the ingress port is not a member of the classified VLAN, the frame is discarded.

Frame Type

Shows whether the port accepts all frames or only tagged frames. This parameter affects VLAN ingress processing. If the port only accepts tagged frames, untagged frames received on that port are discarded.

Tx Tag

Shows the egress filtering frame status whether tagged or untagged (**Tag_all** or **(Untag_this** or **Untag_all)**).

UVID

Shows the UVID (untagged VLAN ID) for each port. A port's UVID determines the packet's behaviour at the egress side. Frames transmitted from this port are untagged. Each port can be an untagged member of just one VLAN. By default, all ports are an untagged member of VLAN 1.

Conflicts

Shows the status of Conflicts whether one exists or not. When a Volatile VLAN User requests to set VLAN membership or VLAN port configuration, the following conflicts can occur:

Functional Conflicts between features.

Conflicts due to hardware limitation.

Direct conflict between user modules.

Buttons

Auto-refresh: Check this box to enable an automatic refresh of the page at 3 second intervals.

Refresh: Click to refresh the page immediately.

Examples

An example of the VLAN Port Status for "Combined" users is shown below:

The screenshot shows the NET2EDGE web interface for a LIB-4424 Carrier Ethernet Network Interface Device. The left sidebar contains a navigation menu with options like LLDP, Ethernet Services, Performance Monitor, PTP, MAC Table, VLANs, Membership, Ports, MVRP, sFlow, and LLDP. The main content area is titled 'VLAN Port Status for Combined users'. It features a table with columns: Port, Port Type, Ingress Filtering, Frame Type, Port VLAN ID, Tx Tag, Untagged VLAN ID, and Conflicts. The table lists six C-Ports (1-6) with Ingress Filtering checked, Frame Type set to 'All', and Tx Tag set to 'Untag All'. The 'Conflicts' column for all ports is 'No'. Above the table, there are controls for 'Combined' (dropdown), 'Auto-refresh' (checkbox), and a 'Refresh' button.

Port	Port Type	Ingress Filtering	Frame Type	Port VLAN ID	Tx Tag	Untagged VLAN ID	Conflicts
1	C-Port	<input checked="" type="checkbox"/>	All	1	Tag All		No
2	C-Port	<input checked="" type="checkbox"/>	All	1	Untag All		No
3	C-Port	<input checked="" type="checkbox"/>	All	1	Tag All		No
4	C-Port	<input checked="" type="checkbox"/>	All	1	Untag All		No
5	C-Port	<input checked="" type="checkbox"/>	All	1	Untag All		No
6	C-Port	<input checked="" type="checkbox"/>	All	1	Untag All		No

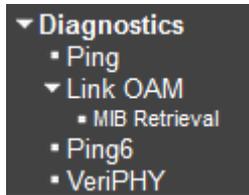
The VLAN Port Status for one “MVR” user is shown below:

The screenshot shows the NET2EDGE web interface for a LIB-4424 Carrier Ethernet Network Interface Device, filtered for an MVR user. The left sidebar is the same as in the previous screenshot. The main content area is titled 'VLAN Port Status for MVR user'. It features a table with columns: Port, Port Type, Ingress Filtering, Frame Type, Port VLAN ID, Tx Tag, Untagged VLAN ID, and Conflicts. The table lists four C-Ports (1-4) with Ingress Filtering checked, Frame Type set to 'All', and Tx Tag set to 'Untag All'. The 'Conflicts' column for all ports is 'No'. Above the table, there are controls for 'MVR' (dropdown), 'Auto-refresh' (checkbox), and a 'Refresh' button.

Port	Port Type	Ingress Filtering	Frame Type	Port VLAN ID	Tx Tag	Untagged VLAN ID	Conflicts
1	C-Port	<input checked="" type="checkbox"/>	All		Tag All		No
2	C-Port	<input checked="" type="checkbox"/>	All		Untag All		No
3	C-Port	<input checked="" type="checkbox"/>	All		Tag All		No
4	C-Port	<input checked="" type="checkbox"/>	All		Untag All		No

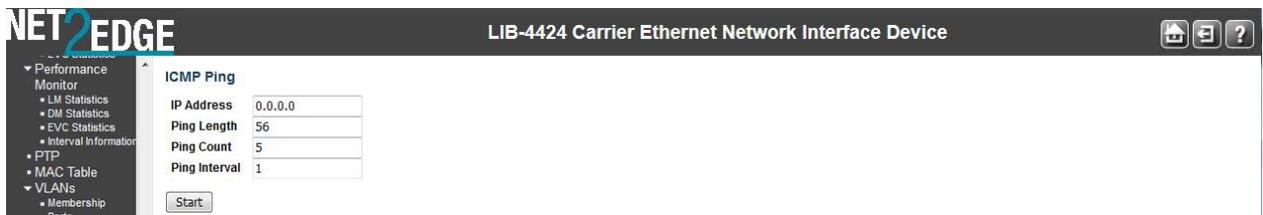
Diagnostics Main Menu

The **Diagnostics** main menu lets you select the **Ping**, **Link OAM**, **Ping6**, **VeriPHY**, and **Service Activation** sub-menus.



Diagnostics > Ping

This page lets you issue ICMP PING packets to troubleshoot IPv4 connectivity issues.



Procedure

Navigate to the **Diagnostics > Ping** menu path.

At **IP Address** enter a valid IPv4 address (e.g., 192.168.1.30).

At **Ping Length** enter the packet size in bytes.

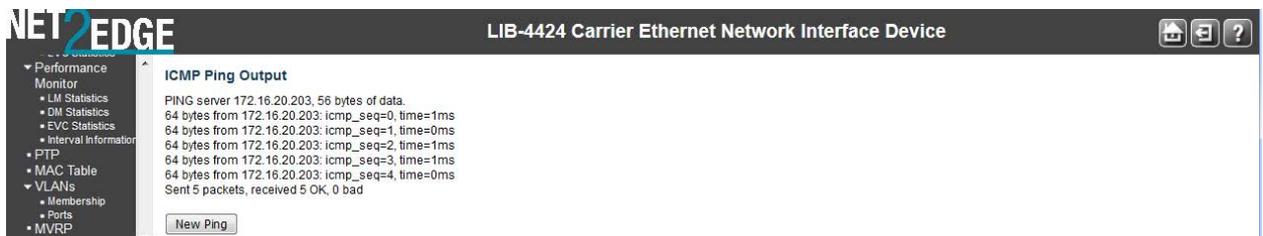
At **Ping Count** enter the number of packets to be sent. The default is 5 pings.

At **Ping Interval** enter the interval to be inserted between pings. The default is 1 msec.

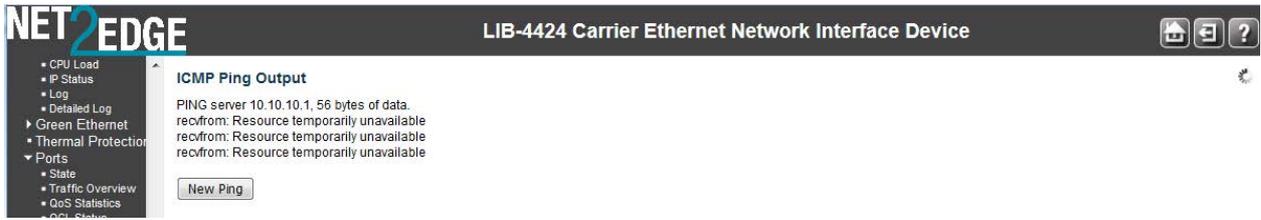
Click on the **Start** button to retrieve the Ping6 output.

After you press the **Start** button, 5 ICMP packets are transmitted, and the sequence number and roundtrip time are displayed upon reception of a reply. The page refreshes automatically until responses to all packets are received, or until a timeout occurs.

A successful ICMP Ping Output is shown below:



A failed ICMP Ping Output is shown below:



Click the **New Ping** button to issue another ping. For example:

```
PING6 server ::10.10.132.20
64 bytes from ::10.10.132.20: icmp_seq=0, time=0ms
64 bytes from ::10.10.132.20: icmp_seq=1, time=0ms
64 bytes from ::10.10.132.20: icmp_seq=2, time=0ms
64 bytes from ::10.10.132.20: icmp_seq=3, time=0ms
64 bytes from ::10.10.132.20: icmp_seq=4, time=0ms
Sent 5 packets, received 5 OK, 0 bad
```

A failed Ping is shown below:

```
ICMP Ping Output
PING server 192.168.1.10
rcvfrom: Operation timed out
Sent 5 packets, received 0 OK, 0 bad
```

Diagnostics > Link OAM > MIB Retrieval

The **Diagnostics > Link OAM > MIB Retrieval** menu path lets you retrieve the Local or Remote OAM MIB variable data on a particular port.



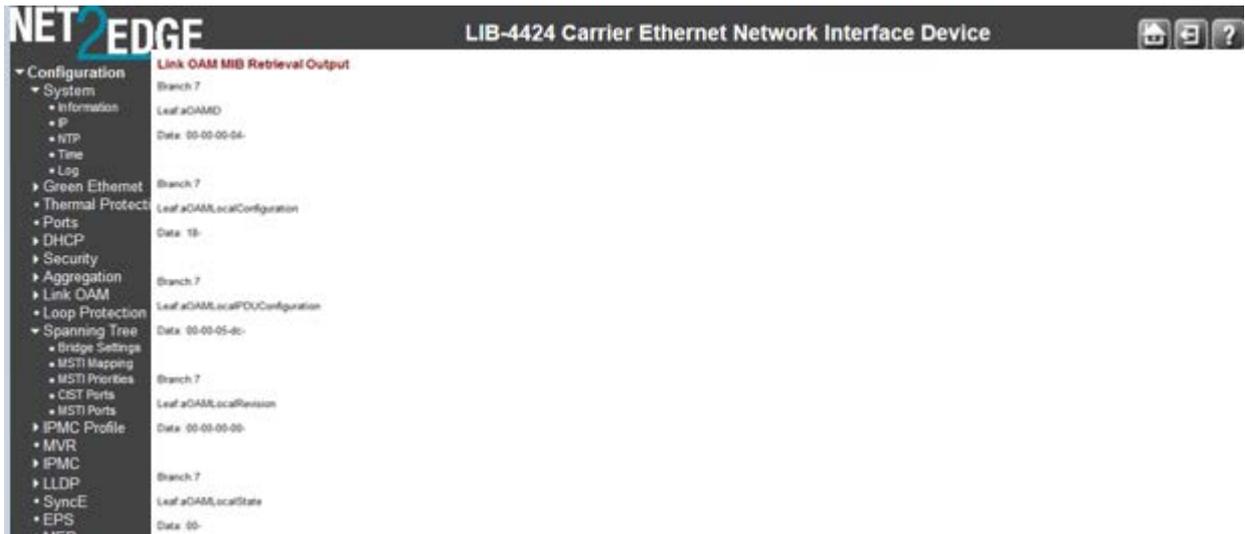
Procedure

Make sure the MIB Retrieval Support checkbox is checked at **Configuration > Link OAM > Port Settings**.

Navigate to the **Diagnostics > Link OAM > Mib Retrieve** menu path.

Select the appropriate radio button to retrieve the content of interest (“Local” or “Peer”). Enter the LIB-44xx Port number (e.g., 1-4). This port must be configured and enabled.

Click the **Start** button to retrieve the MIB content. A typical display is shown below

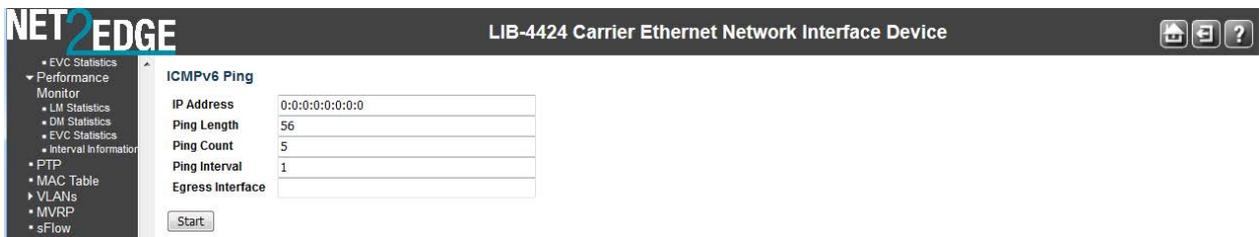


If the Link OAM Mib Retrieve fails, the message “OAM Error - Invalid request on this port” displays. Click the browser’s Back button to clear the message, verify your selections, and then try the operation again.

Note: the **Monitor > Link OAM > Port Status** page provides detailed Link OAM status.

Diagnostics > Ping6

This page lets you issue ICMPv6 PING packets to troubleshoot IPv6 connectivity issues.



Procedure

Navigate to the **Diagnostics > Ping6** menu path.

At **IP Address**, enter a valid IPv6 address. This is the destination IP Address for the ping.

Enter a **Ping Length** (8 - 1400 bytes). This is the payload size of the ICMP packet. The valid values are:

8 to **1400** bytes. The default is **56** bytes.

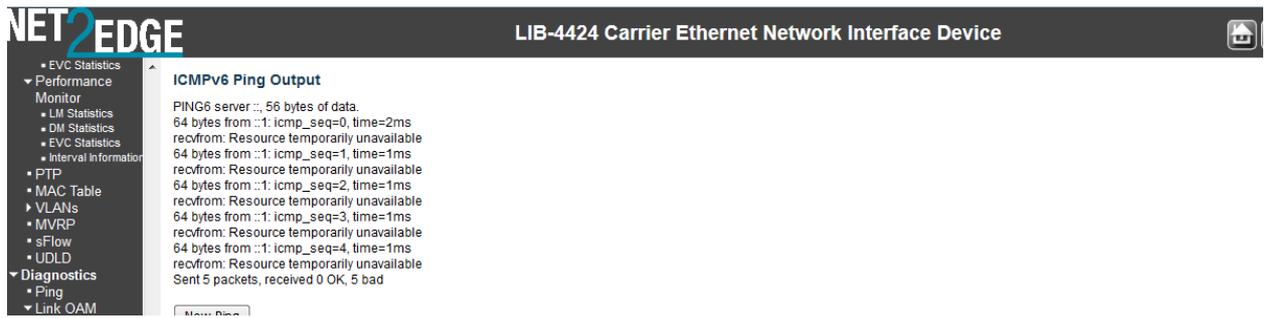
At **Ping Count** enter the number of packets to be sent. This is the count of the ICMP packet. The valid values are **1** to **60** pings. The default is **5** pings.

At **Ping Interval** enter the interval to be inserted between pings. This is the interval of the ICMP packet. Valid values are **0** to **30** seconds. The default is **1** second.

Click on the **Start** button to retrieve the Ping6 output.

After you press the **Start** button, five ICMPv6 packets are transmitted, and the sequence number and roundtrip time are displayed when a reply is received. The page refreshes automatically until responses to all packets are received, or until a timeout occurs.

A successful Ping6 result is shown below:



A failed Ping6 result is shown below:

You can click the **New Ping** button to re-display the initial Ping6 page.

Diagnostics > VeriPHY

This page is used for running the VeriPHY Cable Diagnostics. The RJ45 ports (10/100/1000BaseT) support this cable test using Time Domain Reflectometry (TDR). The TDR method can detect an open, short or normal condition on each of the pairs. If the test result is “normal”, it displays the cable length. This test is intrusive since the port’s link is brought down.

The VeriPHY Cable Diagnostics enable a variety of cable operating conditions and status to be accessed and checked. The VeriPHY suite identifies the cable length and operating conditions, and isolates various common faults that can occur on the LIB-44xx CAT 5 twisted pair (TP) cabling. If a link is established on the TP in 1000BASE-T mode, VeriPHY runs without disrupting the link or disrupting any data transfer. However, if a link is established in 100BASE-T or 10BASE-T, VeriPHY causes the link to drop while the diagnostics are running. When diagnostics are done running, the link is then re-established. These functions are part of the VeriPHY suite:

Detecting coupling between cable pairs: shorted wires, improper termination, or high crosstalk resulting from an incorrect wire map can cause error conditions, such as an anomalous coupling between cable pairs. All of these conditions can prevent the devices from establishing a link at any speed.

Detecting cable pair termination: proper termination of CAT5 cable requires 100 ohms of differential impedance between the positive and negative cable terminals. The IEEE 802.3 standard allows for a termination of 85 - 115 ohms. If the termination falls outside of this range, it is reported by the VeriPHY diagnostics as an anomalous termination. The VeriPHY diagnostics can also determine the presence of an open or shorted cable pair.

Determining cable length: when the CAT5 cable in an installation is properly terminated, VeriPHY reports the approximate cable length in meters (m).

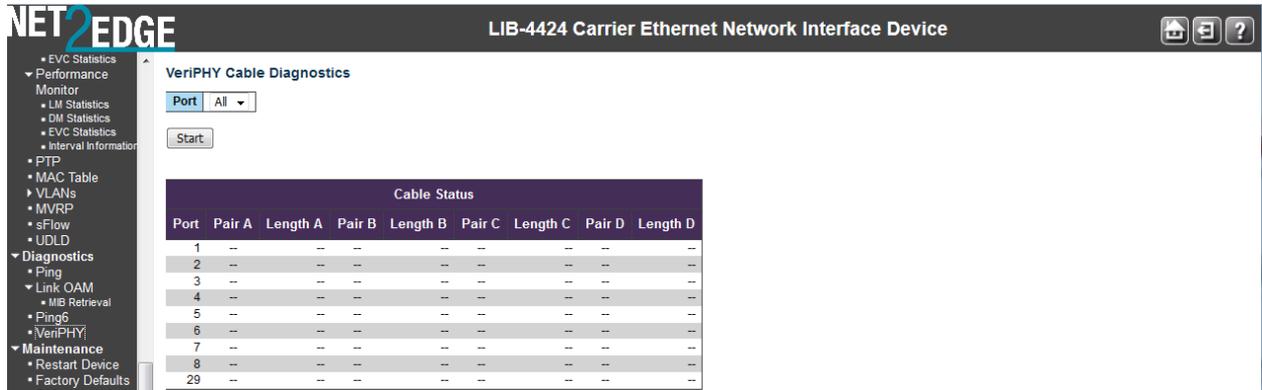
Operating Notes

VeriPHY Multi-Port Operation: On multi-port devices, the VeriPHY cable diagnostics can be executed on a specific port or on all the ports simultaneously.

VeriPHY Interaction with Normal PHY Operation: If a link is established on the twisted pair interface in 1000BASE-T mode, VeriPHY cable diagnostics can run without disruption of the link or of any data transfer. However, if a link is established in 100BASE-TX or 10BASE-T, the VeriPHY cable diagnostics will cause the link to drop while the diagnostics are running. Once the diagnostics are finished, the link will be reestablished. During the time that the function is running, the PHY registers will not be available. This may affect the operation of the link status polling algorithms in some cases.

VeriPHY Range: Shorted wires, improper termination, or high crosstalk resulting from an incorrect wire map can cause anomalous coupling between cable pairs. All of these conditions can prevent the PHY from establishing a link at any speed. The VeriPHY feature can correctly identify a cross-pair short location at up to 100 meters, with 10-meter accuracy.

The default VeriPHY page is shown below:



Procedure

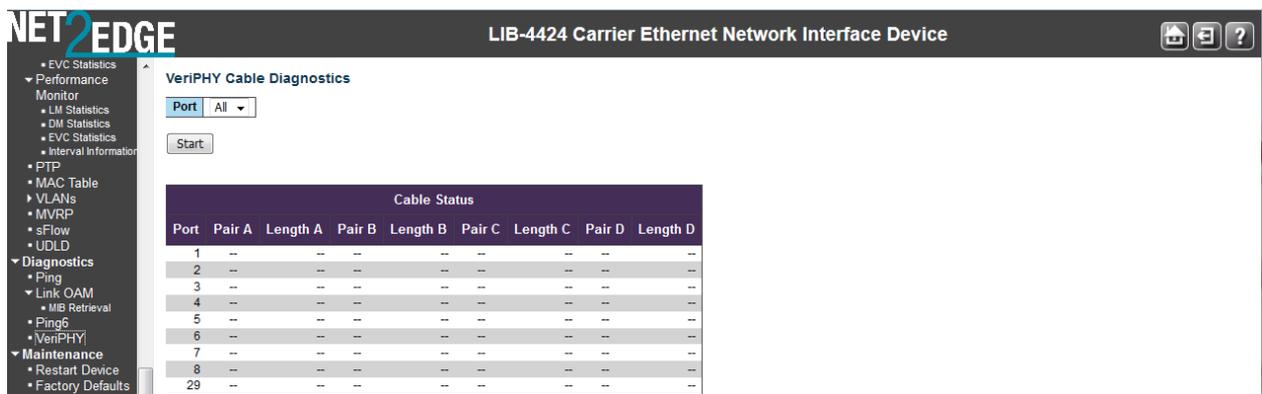
Navigate to the **Diagnostics > VeriPHY** menu path.

Select a specific port (1-4) or **All** ports from the **Port** dropdown. This is the port where you are requesting VeriPHY Cable Diagnostics.

Press the **Start** button to run the diagnostics. The message “x *VeriPHY is running...*” displays momentarily and the  icon displays. The VeriPHY diagnostic takes about five seconds for a single port, or about 15 seconds if ‘All’ ports are selected. When completed, the page refreshes automatically, and you can view the cable diagnostics results in the cable status table.

Note: VeriPHY is only accurate for cable lengths of 7 - 140 meters.

A completed VeriPHY page is shown below:



Note: The LIB-44xx 10 Mbps ports and 100 Mbps ports will be linked down while running VeriPHY. Therefore, running VeriPHY on a 10 Mbps or 100 Mbps management port will cause the LIB-44xx to stop responding until VeriPHY is complete.

Cable Status

Port: The Port number under test.

Pair x: The status of the cable pair (e.g., **Open**, **Short**, **OK**).

Length: The length (in meters) of the cable pair (7 - 140 meters).

Pair Status: The status of the cable pair.

Open: Open pair. The test revealed an open in the cabling for the specified port.

OK: Correctly terminated pair. The test revealed no problems in the cabling for the specified port.

Short: Shorted pair. The test revealed a circuit short in the cabling for the specified port.

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.

Messages

Message: *x VeriPHY is running...*

Meaning: The test is in process.

Recovery:

1. Wait for completion or another message.
2. Click the browser's Back button, and then click the Forward button.
3. Contact Tech Support if the problem persists.

Message: *Switch is currently not responding. Please wait...*

Meaning: The test has encountered a problem.

Recovery:

1. Wait for completion or another message.
2. Click the browser 'Back' button.
3. Switch to another menu path and then switch back to the **Diagnostics > VeriPHY** menu path.
4. Contact Tech Support if the problem persists.

Diagnostics > Service Activation

The **Diagnostics > Service Activation** menu path lets you run a defined SA test. The SA test must first be configured from the **Configuration > Service Activation** menu path.

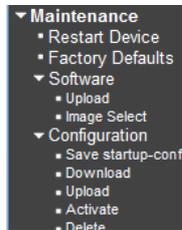
When the System, Profiles and Test pages are configured you can run the SA tests from the **Diagnostics > Service Activation > Test** menu path.

Note: Policy ID 254 is used for marking traffic. Make sure this Policy ID is not used for other purposes (ECs) and the ACE Policy Filter is not being used as a bit field that would inadvertently match 254 (i.e., Policy Bitmask should be 0xFF for all ACEs).

See the *EtherSAT User Guide* manual for EtherSAT configuration and operation.

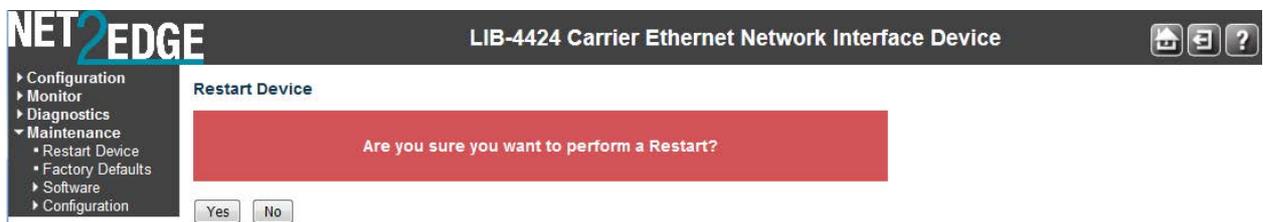
Maintenance Menu

The LIB-44xx **Maintenance** menu provides the **Restart Device**, **Factory Defaults**, **Software** and **Configuration** menus and sub-menus.



Maintenance > Restart Device

You can restart the LIB-44xx from this page. After a restart, the LIB-44xx will boot normally.



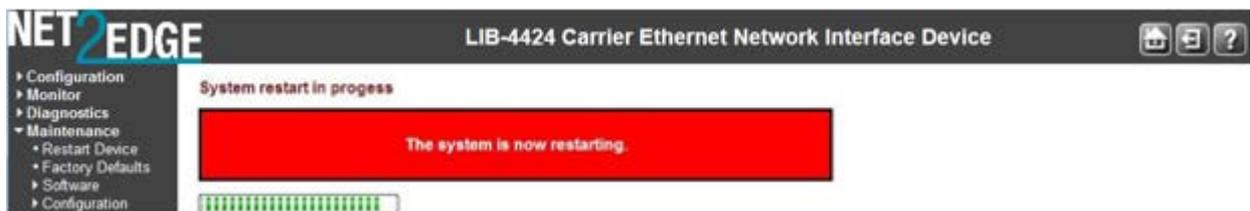
Procedure

Navigate to the **Maintenance > Restart Device** menu path. The confirmation message “*Are you sure you want to perform a Restart?*” displays.

If you are sure you want to restart the LIB-44xx, click the **Yes** button.

If you are not sure you want to restart the LIB-44xx, click the **No** button and continue operation.

To restart the LIB-44xx, click the **Yes** button.



The “*System restart in progress*” screen displays with a series of messages, starting with “*Waiting, please stand by ...*”. When the restart is complete, the LIB-44xx startup screen (**Monitor > Ports > State** page) displays.

Buttons

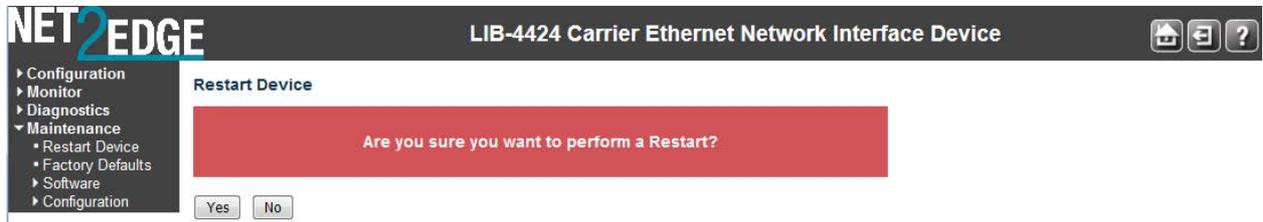
Yes: Click to restart device.

No: Click to return to the Port State page without restarting.

Maintenance > Restart Device > Force Cool Restart

An error condition may display the **Restart Device** page with an option to force an LIB-44xx cool restart.

If this occurs, at **Maintenance > Restart Device > Are you sure you want to perform a Restart?** - **Force Cool Restart**, check or uncheck the checkbox and click the **Yes** button.



Check the **Force Cool Restart** checkbox and click **Yes** to perform an LIB-44xx cool restart.

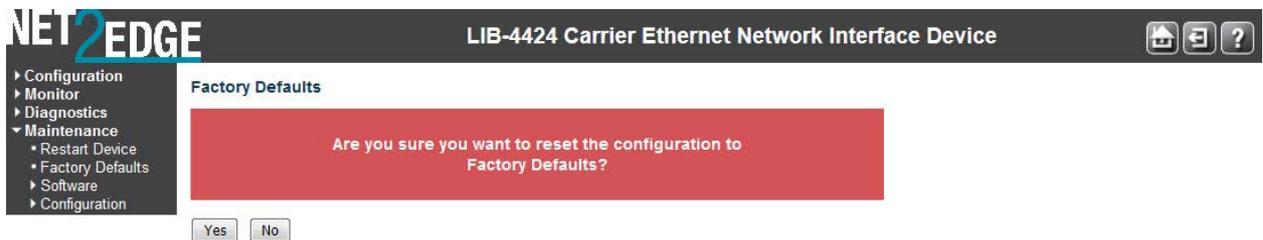
Uncheck the **Force Cool Restart** checkbox and click **Yes**, to perform an LIB-44xx cool restart.

Click the **No** button to clear the message without performing any restart.

Maintenance > Factory Defaults

You can reset the LIB-44xx configuration to its factory default settings from this page.

See “[Appendix C - Application Notes](#)” on page 511 for the full set of LIB-44xx factory default settings.



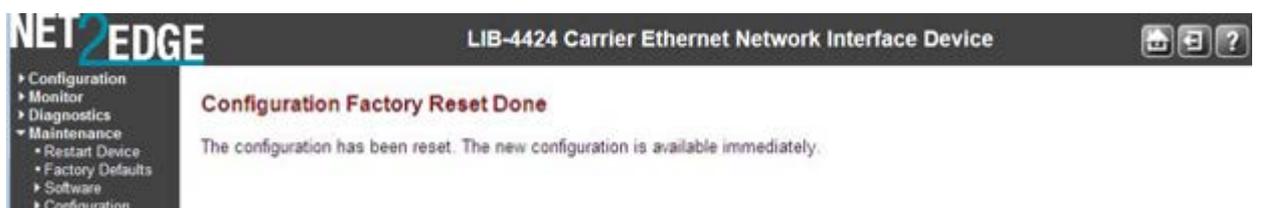
Only the IP configuration is retained after a reset to factory configuration is performed. The new configuration is available immediately, which means that no restart is needed.

Procedure

Navigate to the **Maintenance > Factory Defaults** menu path. The confirmation message “*Are you sure you want to reset the configuration to Factory Defaults?*” displays.

If you are not sure you want to restart the LIB-44xx, click the **No** button and continue operation.

If you are are sure you want to restart the LIB-44xx, click the **Yes** button. The information message “*Configuration Factory Reset Done - The configuration has been reset. The new configuration is available immediately.*” displays.



Continue operation.

Buttons

Yes: Click to reset the configuration to Factory Defaults.

No: Click to return to the Port State page without resetting the configuration.

Maintenance > Software

The LIB-44xx **Maintenance** > **Software** menu path lets you select the **Upload** and **Image Select** sub-menus.

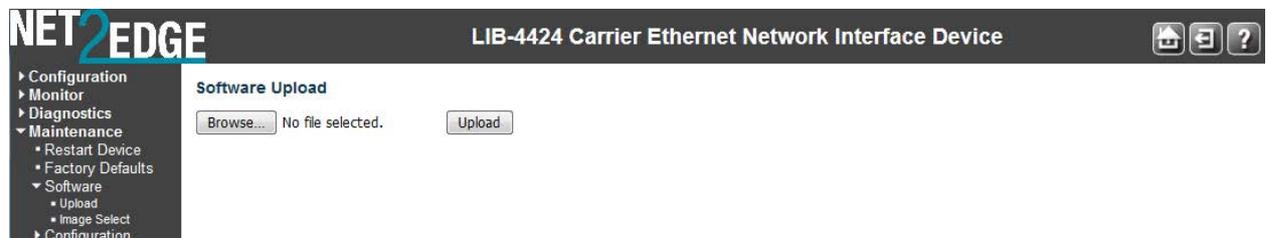
Software Upload via the Maintenance > Software > Upload Path

The LIB-44xx supports firmware upgrade via TFTP and HTTP from the **Maintenance** > **Software** > **Upload** path (the Web uses HTTP, and the CLI uses TFTP).

All configuration settings are retained when the device resets after successful upgrade operation. If a previous configuration is not compatible with the new, it will be set to factory default. Specific messages can be viewed in the CLI during reboot after the update is complete. The firmware image has a CRC mechanism to prevent a corrupted image from being loaded onto the LIB-44xx. The upgrade procedure handles error conditions such as a network disruption during upgrade, power outages, TFTP/HTTP connection issues, etc. and will continue to operate using the installed image in case of upgrade failures.

It is a good idea to create a backup of the configuration before upgrading the firmware.

The LIB-44xx does not disrupt any data plane traffic during the image download process; the data traffic will experience a loss when the device resets to boot with the new image, but the service is restored immediately after the LIB-44xx is configured. The *.dat* file contains a checksum which is validated after upload. If power outage occurs while writing the flash, the flash upgrade will fail. This is why there is an alternate image to serve as an alternate or replacement file.



This page lets you update the LIB-44xx firmware.

Warning: While the firmware is being updated, Web access appears to be defunct. The front panel **S1** 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 LIB-44xx may fail to function afterwards. The upload sequence cannot 'continue' after a disruption. A new upload sequence must be initiated.

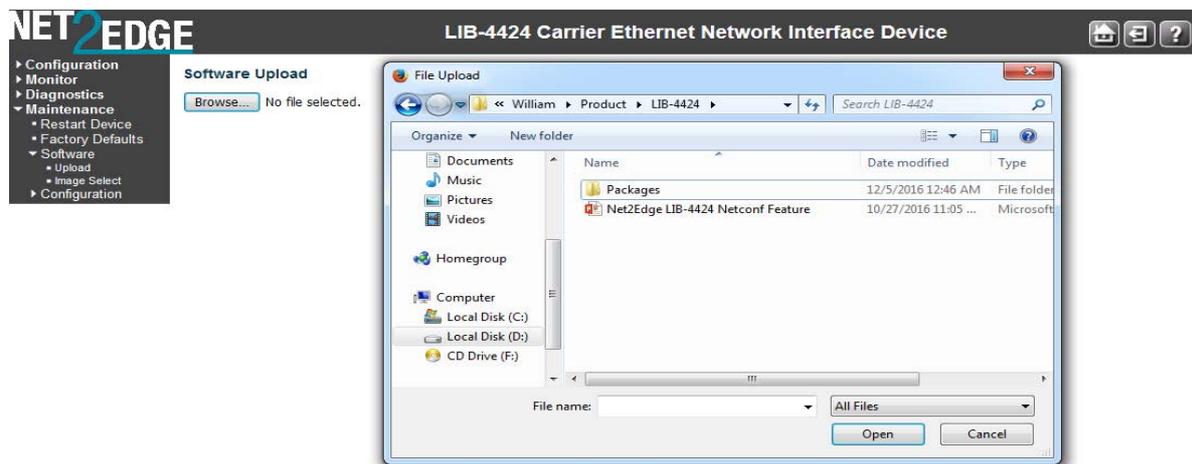
Software Upload Procedure

This procedure transfers the uploaded firmware image to the LIB-44xx flash component. You have the option to activate the image immediately or later.

Note: do not reset or power off the LIB-44xx during this process.

Navigate to the **Maintenance > Software > Upload** menu path.

Click the **Browse** button. The “Choose File to Upload” dialog displays.



Browse to the location of a software image, select a file name with a **.DAT File** extension, and click the **Open** button.

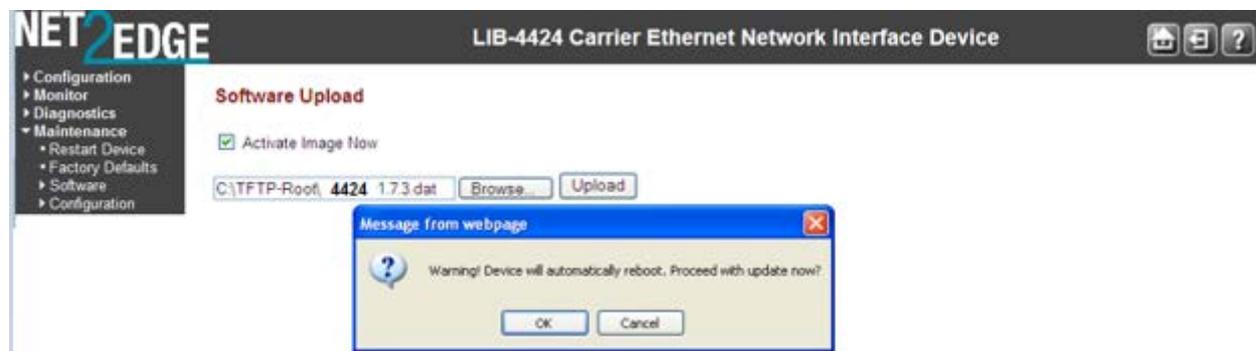
At the **Activate Image Now** checkbox, check or uncheck the box:

If you leave the checkbox unchecked, the image will be uploaded, but not immediately activated.

If you check the checkbox, the image will be uploaded, and can be immediately activated.

Click the **Upload** button. The confirmation dialog "*Warning! Device will automatically reboot. Proceed with update now?*" displays.

If the upload version already is installed, the message "*Firmware Upload Error - Flash is already updated with this image*" displays. Click the browser Back button to recover.

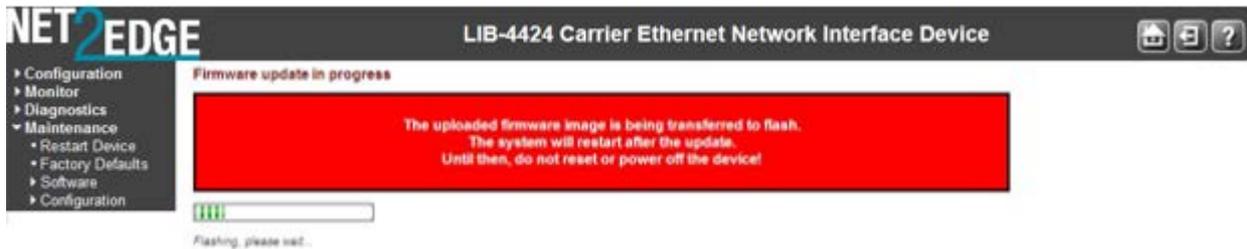


Note: do not reset or power off the LIB-44xx until this firmware update procedure completes. The upload sequence cannot ‘continue’ after a disruption. A new upload sequence must be initiated.

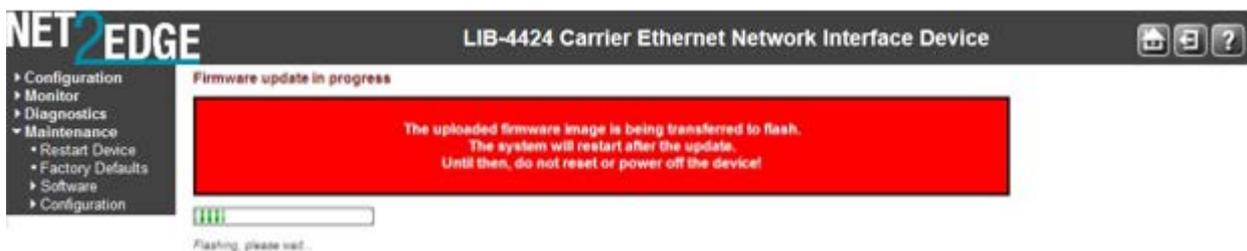
Click **OK** to proceed with the update, or click **Cancel** to return to the Firmware Update page. After the software image is uploaded, a page announces that the firmware update is initiated.

Note: do not reset or power off the LIB-44xx during this process. The upload sequence cannot ‘continue’ after a disruption. A new upload sequence must be initiated.

If you checked the **Activate Image Now** checkbox, the message “*The upload firmware image is being transferred to flash. The system will restart after the update. Until then, do not reset or power off the device!*” displays.



If you unchecked the **Activate Image Now** checkbox, the message “*Writing firmware image to flash. Do not reset or power off the device!*” displays.



The “*Flashing, please wait ...*” series of messages display during the process. After 1-2 minutes, the firmware is updated.

If you checked the **Activate Image Now** checkbox, the LIB-44xx restarts. When the LIB-44xx startup screen displays, continue operation.

If you unchecked the **Activate Image Now** checkbox, the **Software Image Selection** page displays. Continue with the **Image Select** procedure below:

Messages:

Do not reset or power off the device!

Error: Incomplete stack update - update aborted

FIRMWARE_ERROR_xxx code

Flashing, please wait...

Flash is already updated with this image

Programming, please wait ...

Rebooting system...

Restarting, please wait...

Slave, only doing local update

The uploaded firmware image is invalid. Please use a correct firmware image.

Waiting for firmware update to complete

(Still) waiting for firmware update to complete

Warning! Device will automatically reboot. Proceed with update now?

Maintenance > Software > Image Select

This page provides information about the Active (current) and Alternate (backup) firmware images in the device, and allows you to revert to the Alternate Image.

NET2EDGE LIB-4424 Carrier Ethernet Network Interface Device

Configuration
Monitor
Diagnostics
Maintenance
 Restart Device
 Factory Defaults
 Software
 Upload
 Image Select
 Configuration

Software Image Selection

Active Image

Image	net2edge_ce_jr2_24.mfi
Version	LIB-4424dev-build by jon@jon-VirtualBox 2016-11-17T14:44:45+00:00 Config:net2edge_ce_jr2_24 SDK:v02.25-smb
Date	2016-11-17T14:44:45+00:00

Alternate Image

Image	linux.bk
Version	dev-build by jon@jon-VirtualBox 2016-11-17T15:42:29+00:00 Config:net2edge_ce_jr2_24 SDK:v02.25-smb
Date	2016-11-17T15:42:29+00:00

Activate Alternate Image Cancel

The web page displays two tables with information about the **Active Image** and the **Alternate Image**.

Note: If the Active Image firmware image is the same as the Alternate image, only the "Active Image" table displays. In this case, the **Activate Alternate Image** button is also disabled.

If the Alternate Image is active (due to a corruption of the primary image or due to manual intervention), uploading a new firmware image to the device will automatically use the primary image slot and activate this.

Image

The flash index name of the firmware image. The name of the active (existing) image is **managed**; the alternate image is named **managed.bk**.

Version

The version of the firmware image (e.g., as shown above, **LIB-4400 (standalone) 1.7.3**).

Date

The date when the firmware was produced (e.g., **2013-08-20T21:05:48-05:00** as shown above).

Image Select Procedure (Activate Alternate Image)

Navigate to the **Maintenance > Software > Upload** menu path.

NET2EDGE LIB-4424 Carrier Ethernet Network Interface Device

Configuration
Monitor
Diagnostics
Maintenance
 Restart Device
 Factory Defaults
 Software
 Upload
 Image Select
 Configuration

Software Image Selection

Active Image

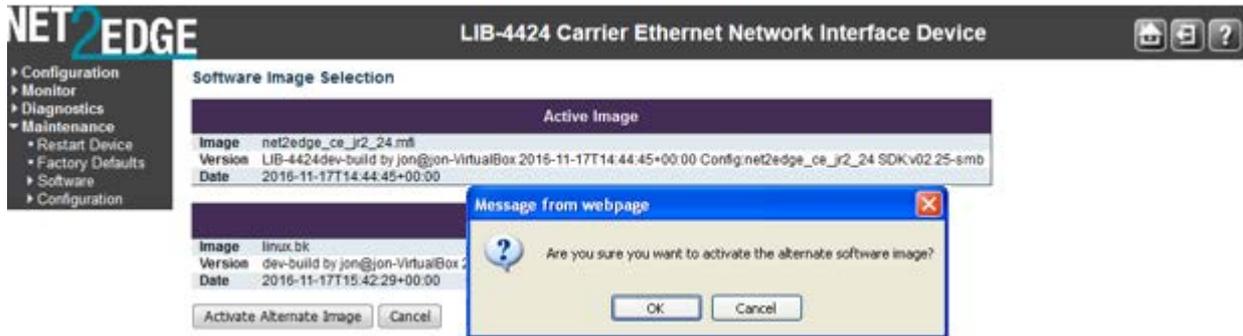
Image	net2edge_ce_jr2_24.mfi
Version	LIB-4424dev-build by jon@jon-VirtualBox 2016-11-17T14:44:45+00:00 Config:net2edge_ce_jr2_24 SDK:v02.25-smb
Date	2016-11-17T14:44:45+00:00

Alternate Image

Image	linux.bk
Version	dev-build by jon@jon-VirtualBox 2016-11-17T15:42:29+00:00 Config:net2edge_ce_jr2_24 SDK:v02.25-smb
Date	2016-11-17T15:42:29+00:00

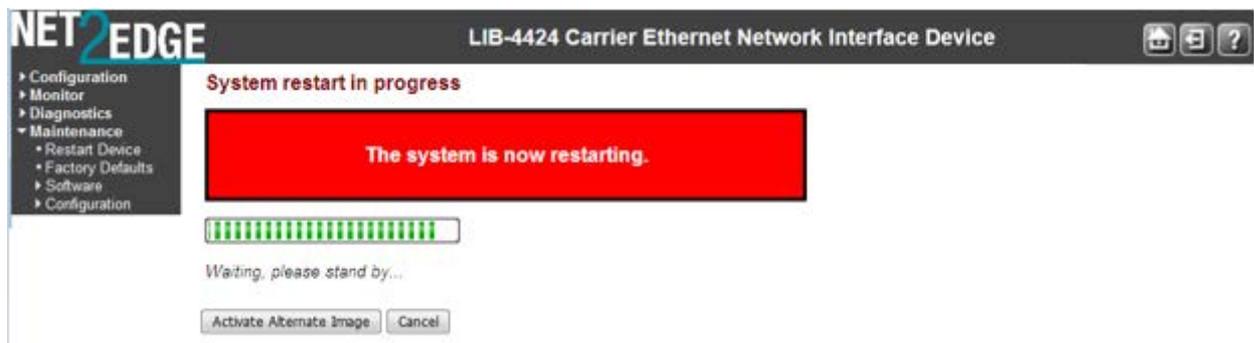
Activate Alternate Image Cancel

Click the **Activate Alternate Image** button. A confirmation message displays:



If you are not sure you want to activate the alternate LIB-44xx image, click the **Cancel** button and continue operation.

If you are sure you want to restart the LIB-44xx, click the **OK** button. The LIB-44xx restarts and displays the “System restart in progress” message as shown below:



When the LIB-44xx startup screen (**Monitor > Ports > State**) displays, continue operation.

Buttons

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.

Messages:

- Activate image (swap) now and reboot*
- Activate image (swap) manually later*
- Alternate image activated, now rebooting.*
- Alternate image activation failed.*

Peripheral Device Firmware Update Messages

Message:

Device not present

FPGA Version: Not present

FPGA Version: Unknown

Meaning: No peripheral device firmware to update.

Recovery:

1. Verify that your LIB-44xx model supports this feature.

2. Try a different function on this device.
3. Try this function on a different device.

Message:

Device flash update complete...\n

Meaning: Successfully completed Peripheral Device Firmware Update procedure.

Recovery:

1. Continue operation.

Message:

Do not reset or power off the device!

Erasing device, please wait...

Erasing device image...\n

Meaning: Peripheral Device Firmware Update procedure is in process.

Recovery:

1. Wait for the upload to successfully complete.
2. Continue operation.

Message: ET-PIC is not present

Meaning: The ET-PIC is not present if the et_board_id does not include timing.

Recovery:

1. Verify that your LIB-44xx model supports this function.
2. Try a different function on this device.
3. Try this function on a different device.

Message:

Peripheral Device Firmware Error - firmware file is too short

Error: cksum mismatch

Error: Firmware file is too short

Error: Firmware file contains bad length

Error: Firmware update failed

esce_id: %lu not found", esce_id

FPGA firmware update failed

FPGA Version: v%u.%u\n,(version.major, version.minor)

HPIC control data not present

HPIC firmware update failed

HPIC usb data not present

illegal esce id: %lu", esce->id

Must explicitly add case for module ID: %d

Opening device failed\n

Meaning: Peripheral Device Firmware Update procedure was unsuccessful.

Recovery:

1. Wait for the upload to successfully complete.
2. Verify the “[Peripheral Device Firmware Update Procedure](#)” steps and entries above (page [Error! Bookmark not defined.](#)).
3. Continue to troubleshoot the peripheral device firmware update / upload problem.
4. Continue operation.

Message:

Peripheral Device Firmware Error

Peripheral device firmware update failed
 Peripheral device firmware update in progress
 Programming device, please wait...
 Programming device image...\n
 Programming device complete, please wait...
 Programming device complete...\n
 Reading version failed\n
 Restarting device...\n
 Restarting device, please wait...

Starting device flash update - do not power off device!
 Writing peripheral device firmware image to flash.

Meaning: Peripheral Device Firmware Update procedure is continually “in process” or has failed.

Recovery:

1. Wait for the upload to successfully complete.
2. Verify the “[Peripheral Device Firmware Update Procedure](#)” steps and parameter entries above (page [Error! Bookmark not defined.](#)).
3. Continue to troubleshoot the peripheral device firmware update / upload problem.
4. Continue operation.

Message:

User Configuration area is locked for address 0x%08X

User Configuration device id of 0x%02X does not match image of 0x%02X

User Configuration hw type of 0x%02X does not match image of 0x%02X

* The exception to the rule case: S3820-TST allows for sw function upgrades of 1 (MACSwap Only) and/or 2 (MACSwap & 2544) in same user space.

User Configuration sw function of 0x%02X does not match image of 0x%02X

User Configuration sw function of 0x%02X is locked for image 0x%02X

Meaning: A peripheral device firmware upgrade error occurred.

Recovery:

1. Verify the “[Peripheral Device Firmware Update Procedure](#)” steps and parameter entries above (page [Error! Bookmark not defined.](#)).
2. Continue to troubleshoot the peripheral device firmware update / upload problem.
3. Continue operation.

Message:

Altera SPI Flash status Error: 0x%08X

Altera SPI Flash status timeout error: reg=0x%08X, status_bits=0x%08X

CRC error during application configuration

External configuration reset (nCONFIG) assertion

FIFO contains %d bytes, but expected 4

FPGA Running in unknown/undefined mode

nSTATUS asserted by an external device as the result of an error

User Watchdog Timer timeout

Meaning: A peripheral device firmware upgrade error occurred.

Recovery:

1. Verify the “[Peripheral Device Firmware Update Procedure](#)” steps and parameter entries above (page [Error! Bookmark not defined.](#)).
2. Continue to troubleshoot the peripheral device firmware update / upload problem.
3. Continue operation.

ET-FPGA Version Compatibility Notes

The Main Firmware is pinned to the major release number of FPGA Firmware. So at a specific release of Main Firmware (e.g., v1.3.3) the required FPGA firmware should begin with major version “1” (e.g., v1.0, v1.1, etc.). For example:

Main Firmware Version: v1.3.3

ET-FPGA Firmware Version: v1.x

This means that the ET-FPGA firmware can be released independently as long as the major version of ET-FPGA Firmware does not change. When the ET-FPGA major version is bumped, a new (compatible) version of Main Firmware will be released.

The LIB-44xx provides an HPIC Auto firmware Sync capability where the HPIC is automatically upgraded/downgraded based on the firmware version stored in main firmware.

3. Messages and Troubleshooting

This section provides general and specific LIB-44xx problem solving suggestions, general error recovery steps, and specific web interface messages, meanings, examples, and possible recovery steps.

LIB-44xx Troubleshooting

Check the LIB-44xx [Back Panel Connections](#) (see page 30).

Verify the Installation. Check the Operating System, Web Browser, Telnet Client, and/or Terminal Emulation package support (see the LIB-4400/LIB-4424/ Install Guide).

Make sure your particular model supports the function attempted.

Check the LIB-44xx Front Panel Connectors [and LEDs](#) (see the /LIB-4424/ Install Guide).

Respond to any LIB-44xx error messages (see “LIB-44xx Error Recovery” below).

Run the LIB-44xx Diagnostics tests and verification functions (e.g., Ping, Link OAM Mib Retrieve, Ping6, VeriPHY). See the “[Diagnostics Main Menu](#)” section on page 411.

Perform the LIB-44xx troubleshooting and service functions (e.g., Restart Device, reset to Factory Defaults, Software Upload, Image Select). See the “[Maintenance Menu](#)” section on page 417.

Check the LIB-44xx operating parameters (e.g., Information, CPU Load, Log, Detailed Log). See the “[Monitor](#)” section on page 310.

If you can access the LIB-44xx via PuTTY or HyperTerminal but not via the web interface, enter the **restore default keep ip** CLI command and try accessing the LIB-44xx web interface again.

Ethernet SAT (Service Activation Testing)

The MEF SAT (Service Activation Testing) is implemented early in the Ethernet Service lifecycle; when a new customer order is received, MEF SAT (along with MEF LLB and ITU Y.1564) can be used to provision and turn up the circuit in order to verify the performance to the SLA (via FM and PM). Ethernet Service Activation Test (EtherSAT) methodology involves:

Verify a new service after provisioning is complete, but before it is turned over to the customer. Check that the configuration is correct.

Verify performance meets the Service Acceptance Criteria (SAC) to ensure CoS Performance Objectives are attained.

The EtherSAT loopback test can be run via the Web interface or the CLI. See the related manual for detailed information.

Sync-E and PTP Troubleshooting

Selecting SyncE or IEEE 1588v2 depends on whether the mobile technology application needs frequency or frequency and phase. SyncE is a physical layer technology that delivers a frequency reference, while IEEE 1588v2 delivers time from which both phase and frequency can be derived. There are some caveats:

SyncE cannot be deployed on legacy Ethernet networks unless the physical hardware or interfaces are all upgraded. This may limit its deployment across operator boundaries or national boundaries. On the plus side, it is not affected by network impairments such as frame delay range (FDR).

IEEE 1588v2 can be deployed with legacy Ethernet NEs not implementing the protocol, but clock recovery algorithms at slave clocks need to function well in the presence of network impairments that can vary significantly with traffic load.

Synchronization is essential to support frequency accuracy and stability. Lack of stability or accuracy will cause bit errors and/or underflows and overflows of frame buffers that result in lost packets in the PDH framing, severely affecting traffic performance.

Sync-E Troubleshooting

Many services that are provided over networks must be fully synchronized in order to operate correctly. If the network devices included in the network do not operate at the same clock rates, an overall decrease in the performance of the network occurs, along with a consequent degradation in the quality of service offered by the network.

Note: Your particular LIB-44xx software version may not support all of the features documented in this chapter. For the latest feature details and caveats, see the Release Notes for your particular model and software version.

Problem: *Clock selection*

Solution:

1. Verify that there are no alarms on the interfaces. Use the **Configuration > SyncE** menu path to check this.
2. Make sure that the nonrevertive configurations are in place.
3. See “[Clock Selection Mode and State](#)” on page 288 for more information.

Problem: *Incorrect QL values*

Solution:

1. Make sure that no framing mismatch exists with the SSM option.
2. See “[SyncE \(Synchronous Ethernet\)](#)” on page 286 for more information.

Problem: *Incorrect clock limit set or queue limit disabled mode*

Solution:

1. Verify that no alarms exist on the interface(s). Use the **Monitor > SyncE** menu path to verify.
2. Use the **Monitor > SyncE** menu path to confirm if the system is in revertive mode or nonrevertive mode and verify the nonrevertive configurations.
3. See “[SyncE \(Synchronous Ethernet\)](#)” on page 286 for more information.

IEEE 1588 (PTP) Troubleshooting

A disturbance pattern, as input to the test equipment, could be caused by an FDR profile from a file or an algorithm, a fixed latency, the introduction of errored packets, repeated packets, or any combination of these effects. These effects will impair the flow of interest as it passes from the IEEE 1588v2 master to the slave under test.

Packet impairments are lost packets, mis-ordered packets, re-ordered packets, etc. Performance testing can include packet impairments applied to IEEE 1588v2 timing transfer packets by a network impairment device.

A wander measurement instrument is used to compare the original clock time signature to a time signature that has passed through the network for applications that transfer frequency. Time of day error is measured for applications that use IEEE 1588v2 to transfer time.

During deployment installation, the operator performs turn-up testing on each link to ensure accurate synchronization is being delivered. Once the SDA is verified, tests and system performance monitoring are done on a routine and scheduled basis to assure high quality and that service level agreements are met.

With IEEE 1588v2, the operator is currently required to take a test-and-deploy strategy to determine the SDA most suited to their requirements. In the future this strategy will evolve, once additional rules are defined by relevant specifications.

EPS Troubleshooting

Provisioning Mismatches

With all of the options for provisioning of protection groups, there are opportunities for mismatches between the provisioning at the two ends. These provisioning mismatches take one of several forms:

Mismatches where proper operation is not possible.

Mismatches where one or both sides can adapt their operation to provide a degree of interworking in spite of the mismatch.

Mismatches that do not prevent interworking.

Not all provisioning mismatches can be conveyed and detected by information passed through the APS communication. There are too many combinations of valid entity numbers to easily provide full visibility of all of the configuration options.

Generally, selecting revertive / non-revertive operation is the same at both ends of the protection group. However, a mismatch of the revertive / non-revertive parameter does not prevent interworking.

See Recommendation ITU-T G.8031/Y.1342 for specifics on linear protection switching for Ethernet Virtual Local Area Network (VLAN) signals.

Request State Priorities

<u>Request / State</u>	<u>Priority</u>
1111 Lockout of protection (LO)	
1110 Signal fail for protection (SF-P)	Highest

1101 Forced switch (FS)	^
1011 Signal fail for working (SF)	
1001 Signal degrade (SD) (Note)	
0111 Manual switch (MS)	
0110 Manual switch to working (MS-W)	
0101 Wait to restore (WTR)	
0100 Exercise (EXER)	
0010 Reverse request (RR)	
0001 Do not revert (DNR)	v
0000 No request (NR)	Lowest

Note: SF-P (Signal fail on the protection transport entity) is a higher priority than any defect that would cause a normal traffic signal to be selected from protection.

Protection Types

The valid protection types are:

000x 1+1 Unidirectional, no APS communication

100x 1+1 Unidirectional w/APS communication

101x 1+1 Bidirectional w/APS communication

111x 1:1 Bidirectional w/APS communication

The values are chosen such that the default value (all zeros) matches the only type of protection that can operate without APS (1+1 unidirectional).

Note that 010x, 001x and 011x are invalid since 1:1 and bidirectional require an APS communication. If the "B" bit mismatches, the selector is released since 1:1 and 1+1 are incompatible, resulting in a defect.

If the "B" bit matches:

If the "A" bit mismatches, the side expecting APS will fall back to 1+1 unidirectional switching without APS communication.

If the "D" bit mismatches, the bidirectional side will fall back to unidirectional switching.

If the "R" bit mismatches, one side will clear switches to "WTR"

Failure of Protocol Defects

The "Failure of protocol" situations for protection types requiring APS include:

Fully incompatible provisioning (the "B" bit mismatch

Working/protection configuration mismatch.

Lack of response to a bridge request (i.e., no match in sent "requested signal" and received "requested signal") for >50ms.

Fully incompatible provisionings and working/protection configuration mismatches are detected by receiving a single APS frame.

Detecting and clearing "failure of protocol" defects are defined in ITU-T G.8021.

Any received 'unknown request' or any 'request for an invalid signal number' is ignored.

ERPS Troubleshooting

Failure of protocol defect: due to errors in provisioning, the ERP control process may detect a combination of conditions which should not occur during "normal" conditions. To warn the operator of such an event, a failure of protocol – provisioning mismatch (FOP-PM) is defined. The FOP-PM defect, detected if the RPL owner node receives one or more No Request R-APS message(s) with the RPL Blocked status flag set (NR, RB), and a node ID that differs from its own. The ERP control process must notify the equipment fault management process when it detects such a defect

condition, and continue its operation as well as possible. This is only an overview of the defect condition. The associated defect and its details are defined in ITU-T G.8021 as amended by its Amendments 1 and 2.

IPv6 Troubleshooting

Start by using these third party resources when performing general IPv6 problem solving:

The standard Windows 7 command-line tools with full IPv6 functionality (Ping, Ipconfig, Pathping, Tracert, Netstat, and Route all support IPv6).

The IPv6-specific tools in the Netsh command.

Address Resolution in Windows 7

In unicast global IPv6 (equal to IPv4 Public) addresses, the 64-bit host portion of the address is derived from the MAC address of the network adapter. The Neighbour Discovery (ND) protocol resolves IPv6 addresses to MAC addresses. The resolution of host names to IPv6 addresses is done by DNS with the exception of link-local (equivalent to IPv4 APIPA) addresses, which resolve automatically. DNS handles records for IPv6 host names similar to IPv4 and also uses pointer (PTR) records to perform reverse lookups. Where DNS is not implemented (e.g., peer-to-peer environments) the Peer Name Resolution Protocol (PNRP) provides dynamic name registration and name resolution.

Verify IPv6 Configuration in Windows 7

The main tool is Ipconfig. The command **ipconfig /all** displays both IPv4 and IPv6 configuration. To display the configuration of only the IPv6 interfaces use netsh. The **netsh interface ipv6 show address** command displays each interface IPv6 address including the interface ID after the % character (the configuration can be accessed via the GUI).

Verify IPv6 Connectivity

ping the local address. Note that if pinging link-local addresses from one host to another, you must include the destination adapter interface ID (e.g., ping fe80::38e7:3df1:f5ff:fd0%13). When pinging site-local (equal to IPv4 Private) addresses you can add the interface ID to ensure that the address is configured on the desired interface. You must add an 'allow' rule for ICMPv6 traffic to pass through each computer's firewall.

Command examples - third party CLI commands for IPv6:

```
ipconfig /all
netsh interface ipv6 show address
ping fe80::38e7:3df1:f5ff:fd0%13)
netsh interface ipv6 delete neighbours
netsh interface ipv6 show neighbours
netsh interface ipv6 delete destinationcache
netsh interface ipv6 show destinationcache
netsh interface ipv6 show route
route print
tracert -d <destination IPv6 address>
pathping -d <destination IPv6 address>
```

For Additional Information

IPv6 Forum at <http://www.ipv6forum.com/>

ARIN (American Registry for Internet Numbers) at https://www.arin.net/knowledge/ipv6_info_center.html

or ARIN wiki at http://www.getipv6.info/index.php/Main_Page

Cisco: <http://www.ciscopress.com/articles/article.asp?p=777892&seqNum=7>

Troubleshooting IPv6 on Windows 7: <http://itexpertvoice.com/home/troubleshooting-ipv6-on-windows-7-and-why-its-worth-the-bother/>

Troubleshooting IPv6 on Windows Servers (Microsoft TechNet): [http://technet.microsoft.com/en-us/library/cc780623\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/cc780623(WS.10).aspx)

IPv6 Auto Config Troubleshooting

Determine whether your particular computer will require reconfiguration. For example, for Microsoft .NET Framework version 2.0 and later, IPv6 is enabled by default. For .NET Framework version 1.1 and earlier, IPv6 is disabled by default. For more information see the MSDN article at <http://msdn.microsoft.com/en-us/library/8db2058t.aspx>. Windows Server 2008 provides complete support for IPv6 and all of its features, and does not need additional installation or configuration. For Windows 7 see <http://technet.microsoft.com/en-us/magazine/2009.07.cableguy.aspx>. For Windows XP see <http://support.microsoft.com/kb/2478747>. For Windows Vista see <http://ipv6.com/articles/general/IPv6-Microsoft-Vista.htm>. For Linux / BSD, see <http://ipv6.com/articles/applications/Linux-and-BSD.htm> or http://tldp.org/HOWTO/html_single/Linux+IPv6-HOWTO/ or your distribution documentation and/or website.

RADIUS Troubleshooting in Windows Server Environments

Microsoft RADIUS implementations differ between Windows Server 2003 and Windows Server 2008.

Windows Server 2003: Internet Authentication Service (IAS) is the Microsoft implementation of a Remote Authentication Dial-in User Service (RADIUS) server and proxy in Windows Server 2003. As a RADIUS server, IAS performs centralized connection authentication, authorization, and accounting for many types of network access, including wireless and virtual private network (VPN) connections.

As a RADIUS proxy, IAS forwards authentication and accounting messages to other RADIUS servers.

In Windows Server 2008, IAS was replaced with Network Policy Server (NPS).

See <http://technet.microsoft.com/en-us/network/bb643123> for more information.

Windows Server 2008: Network Policy Server (NPS) is the Microsoft implementation of a Remote Authentication Dial-in User Service (RADIUS) server and proxy in Windows Server 2008. NPS is the replacement for Internet Authentication Service (IAS) in Windows Server 2003. (NPS is actually more than a replacement for IAS, it does what IAS did and much more.) As a RADIUS server, NPS performs centralized connection authentication, authorization, and accounting for many types of network access, including wireless and virtual private network (VPN) connections. As a RADIUS proxy, NPS forwards authentication and accounting messages to other RADIUS servers. NPS also acts as a health evaluation server for Network Access Protection (NAP).

See <http://technet.microsoft.com/en-us/network/bb629414.aspx> for more information.

Configure FreeRadius or TACACS+ for Correct ADMIN Level

AAA 'keyword attribute'

Problem: When the privilege Levels with the Radius Account to the switch are not sent, you have read-only access. FreeRadius is sending Privilege Level 5 per the default. (This also applies to TACACS+ with *service=shell* and *priv-lvl=x*.)

Meaning: If the LIB-44xx does not see these *attrs* then it defaults to level 1 (minimal access). The LIB-44xx can do vendor specific values of Cisco and Zyxel:

```
Vendor-id: 9 (Cisco) Vendor-type: 1
Vendor-id: 890 (Zyxel) Vendor-type: 3
```

FreeRadius sends a Privilege Level 5 by default. The Keyword or attribute for Transition is vendor_value syntax: "shell:priv-lvl=x" where x is an integer from 0 to 15. For Extreme it is 'Extreme-CLI-Authorization', for ADVA DWDM 'ADVA-ADMIN', for other vendors, something else (e.g., the config for Extreme Switches is *Extreme-CLI-Authorization = 1*).

Recovery:

1. The current security privilege setting for the user must be 15. The LIB-44xx range is 1 to 15 (where 15 is the highest value / fullest possible access to all LIB-44xx functions).
2. See "[AAA Configuration](#)" on page 119 of the Web User Guide for more information on configuring via the web interface. See the "[Security AAA commands](#)" section of the CLI Command Reference for more information on configuring via the CLI. See the **Configuration > Security > Switch > Privilege Levels** menu path. See the "[Security Switch Users](#)" commands section.
3. This works similarly for TACACS+ with *service=shell* and *priv-lvl=x*.

FreeRADIUS includes a RADIUS server, a BSD licensed client library, a PAM library, and an Apache module. The word 'FreeRADIUS' usually refers to the RADIUS server. FreeRADIUS is the most widely deployed RADIUS server in the world, and it is the basis for several commercial offerings. FreeRADIUS supplies the AAA needs of many Fortune-500 companies and Tier 1 ISPs.

FreeRADIUS supports a simple processing language in its configuration files, called "un-language". The goal of the language is to allow simple policies to be written with minimal effort. Those policies are then applied when a request is being processed. Requests are processed through virtual servers (including the default one), in the sections titled "authorize", "authenticate", "post-auth", "preacct", "accounting", "pre-proxy", "post-proxy", and "session". The keywords for the language are a combination of pre-defined keywords and references to loadable module names. Subject to a few limitations, any keyword can appear in any context. The language consists of a series of entries, each one with one line. Each entry begins with a keyword and entries are organized into lists. The language is processed line by line, from the start of the list to the end. Actions are executed per-keyword.

For the FreeRADIUS "RADIUS Attribute List" see <http://freeradius.org/rfc/attributes.html>. See <http://freeradius.org/radiusd/man/unlang.html> for the FreeRADIUS "unlang - FreeRADIUS Processing un-language" page. The FreeRADIUS Version 2 Documentation page is at <http://freeradius.org/doc/>.

TACACS+ (and RADIUS) have generally replaced the earlier protocols in more current networks. TACACS+ uses TCP and RADIUS uses UDP; some administrators recommend TACACS+ because TCP is considered more reliable. While RADIUS combines authentication and authorization in a user profile, TACACS+ separates the two operations. TACACS+ is available from Cisco, shrubby.net, rubyforge.org and others.

Troubleshooting High CPU Load Conditions

After the LIB-44xx completes the boot process, the switch CPU performs two distinct functions simultaneously: it runs the various system processes required for a networked switch, and it sends / receives packets to / from the LIB-44xx hardware. CPU load increases when a system process requires more time or when more network packets are sent and received. Under normal operating conditions, the CPU is busy at least 5 percent of the time.

Since background LIB-44xx processes on its switch timers execute multiple times per second, the LIB-44xx never reports CPU utilization at 0%, even for very simple deployments. Normal data traffic packet switching is done in the LIB-44xx hardware without involving the CPU, so it is not affected by an overly busy CPU.

The CPU becomes too busy when it receives too many packets from the LIB-44xx hardware or when a system process consumes too much CPU time. When either of these functions uses CPU resources to the detriment of the other, the CPU becomes "too busy". For example, if the CPU is receiving numerous packets because of a broadcast storm on the network, it becomes so busy processing all of the packets that other system processes do not have access to CPU resources.

In many instances, high CPU load is normal and does not cause network problems. High CPU utilization becomes a problem when the LIB-44xx fails to perform as expected. CPU utilization spikes caused by a known network event or activity are not problems (even an 85% spike may be acceptable, depending on the cause).

Over time, the switch operates within a certain sustained CPU load range, which is considered the normal operations baseline. You can use the output of the **system load** CLI command or the **Monitor > System > CPU Load** menu path.

The CPU Load percentage is shown at 100 ms, 1 second, and 10 second intervals. All numbers represent running averages. Note that the web interface has an Auto-refresh checkbox to refresh the page automatically every 3 seconds.

Frequent unexplained spikes to the established normal operating baseline, or sudden utilization jumps with no explanations are likely causes for concern.

Below are some common symptoms of high CPU utilization.

High percentages in the CLI command output or web interface graph: check the output of the **system load** CLI command.

Slow performance: services fail to respond (e.g., slow Telnet response or unable to Telnet to the LIB-44xx; slow console response, slow or no ping response).

If you notice any of these symptoms, follow the steps below to alleviate the problem.

Check for a possible security issue. A high CPU utilization can be caused by a security issue, such as a worm or virus in your network. This is especially likely if there have not been recent changes to the network. A configuration change (e.g., adding additional lines to ACLs) can mitigate the effects of this problem.

Collect more information using the show version of the CLI commands (e.g., system config, system log, etc.).

If the LIB-44xx is accessible and you can reproduce the problem, try cycling power to the LIB-44xx. Try lowering or disabling all sys logging. Increase logging buffer size.

Make sure any debug commands are turned off. Contact support for details.

Normal Conditions Causing High CPU Load

A busy CPU is normal in some network deployments. Generally, the larger the Layer 2 or Layer 3 network, the greater the demand on the CPU to process network related traffic. Operations with the potential to cause high CPU utilization can include Spanning Tree, IP Routing table updates, encryption via the LIB-44xx software, fragmentation causing the CPU to reassemble numerous packets, or certain CLI commands (e.g., write memory, show config).

Other events that can cause high CPU utilization may include frequent / large number of IGMP requests, the CPU generating numerous ICMP or traceroute packets, SNMP polling activities, numerous simultaneous DHCP requests (e.g., links being restored to numerous clients), ARP broadcast storms, and/or Ethernet broadcast storms.

For Additional High CPU Load TS Information

For troubleshooting High CPU Utilization in Windows see <http://technet.microsoft.com/en-us/library/bb742546.aspx>.

For troubleshooting High CPU Usage on a Domain Controller see <http://technet.microsoft.com/en-us/library/bb727054.aspx>.

For troubleshooting High CPU Utilization issues using **Tracelog** see <http://blogs.technet.com/b/askperf/archive/2012/01/20/troubleshooting-high-cpu-utilization-issues-using-tracelog-exe.aspx>.

For troubleshooting High CPU Utilization in Linux see the documentation for your particular distribution. The Linux **top** program provides a dynamic real-time view of a running system. It can display system summary information as well as a list of tasks currently being managed by the Linux kernel. The Linux CPU utilization displays CPU stats in the **CPU(s)** row and the **%CPU** column.

A task's share of the elapsed CPU time since the last screen update is expressed as a percentage of total CPU time. The top command produces a frequently-updated list of processes. By default, the processes are ordered by percentage of CPU usage, with only the "top" CPU consumers shown. Type q to exit the top command display when done. You can also install a special package called **sysstat** to take advantage of helpful commands. The **sysstat** package includes system performance tools for Linux (Red Hat Linux / RHEL includes these tools by default).

LIB-44xx Error Recovery

The LIB-44xx displays error and information messages from the CLI and Web interface. This section lists the messages, provides an example, and discusses the message meaning of and possible recovery steps.

As a general troubleshooting step for problems encountered using the LIB-44xx web interface, try the related CLI command. For many messages, recovery involves reviewing the command/function description and verifying the entry selection/syntax. For example, for many CLI messages, the first recovery step would be to refer to the "LIB-44xx CLI Reference" manual.

Generic Message Recovery (e.g., you tried a function, but the operation failed or is still in process):

Wait for a few moments for the operation to complete.

Use the **Help** or **?** command to get assistance (help) on a group of commands or on a specific command.

Make sure this is the function you want and that the device/port/configuration supports this function. Verify the parameters entered and re-try the function. See the related section of this manual for specifics.

Try using the CLI to perform the function. Refer to the “LIB-44xx CLI Reference Guide” manual.

If the “continue **y**(es) **n**(o) prompt” displays, type **y** and press **Enter** to continue.

Use the

Monitor sub-menu functions (System, Ports, Link OAM, MAC Table, VLANS) to view related status, statistics, events, etc. related to a specific function.

Use the [Diagnostics](#) Main Menu sub-menu functions (Ping, Link OAM MIB Retrieval, VeriPHY) to test a general functionality.

Use the [Maintenance Menu](#) sub-menu functions (Restart the LIB-44xx, Reset the LIB-44xx to factory defaults, Upgrade the LIB-44xx firmware).

Specific Messages Recovery:

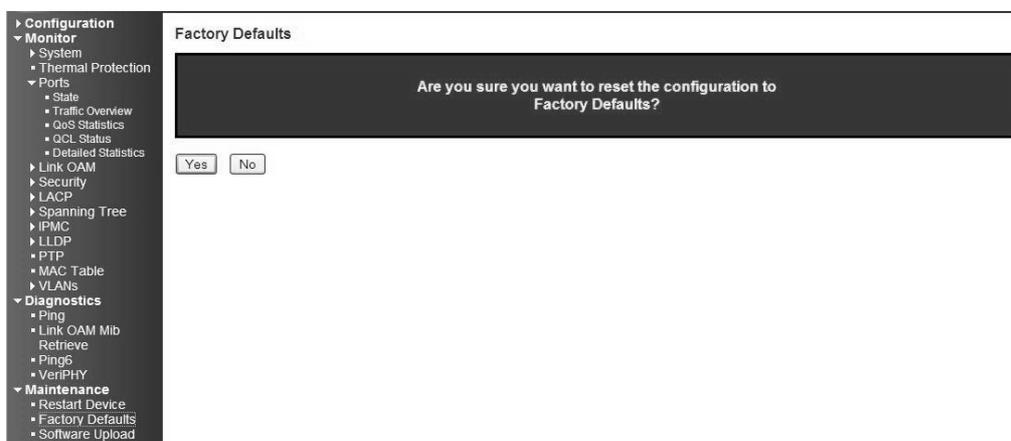
1. For messages (e.g., ACL messages) that are the result of a switch memory shortage:
 - a. Reduce other system activity to ease memory demands.
 - b. Use a less complicated configuration that requires less memory.
 - c. Modify the ACL configuration to use fewer resources, or rename the ACL with a name or number that alphanumerically precedes the other ACL names or numbers.
 - d Reduce the number of IP or MAC access lists to be applied to interfaces.
 - e. Reduce other system activity to ease memory demands (e.g., remove ACLs that are defined but not used; use simpler ACLs with fewer ACEs; use fewer VLANs / remove unneeded VLANs from the VLAN database).
2. For messages that indicate the configuration is too complicated for the ACL code to support, there is likely too many separate access lists in a single VLAN map or policy map. Reduce the number of IP or MAC access lists separately) in any one VLAN or policy map to fewer than the number of levels. Or try to use the same ACLs on multiple interfaces if possible.
3. For messages that indicate an illegal configuration, reconfigure the port / device, removing the illegal configuration.
4. For messages that indicate the temperature is high reduce the temperature in the room.
5. For messages that indicate that the number of MAC address entries for the VLAN exceeds the maximum number allowed, have your system administrator configure an action.
6. For messages that indicate that an unauthorized device attempted to connect on a secure port, identify the device that attempted to connect on the secure port and notify your network system administrator of the condition.

7. For messages that indicate that the amount of traffic detected on the interface has exceeded the configured threshold values, determine and fix the root cause of the excessive traffic on the interface.

8. For messages that indicate an unrecoverable software error has occurred, copy the message exactly as it appears on the console or in the system log and contact Support.

Web Interface Messages

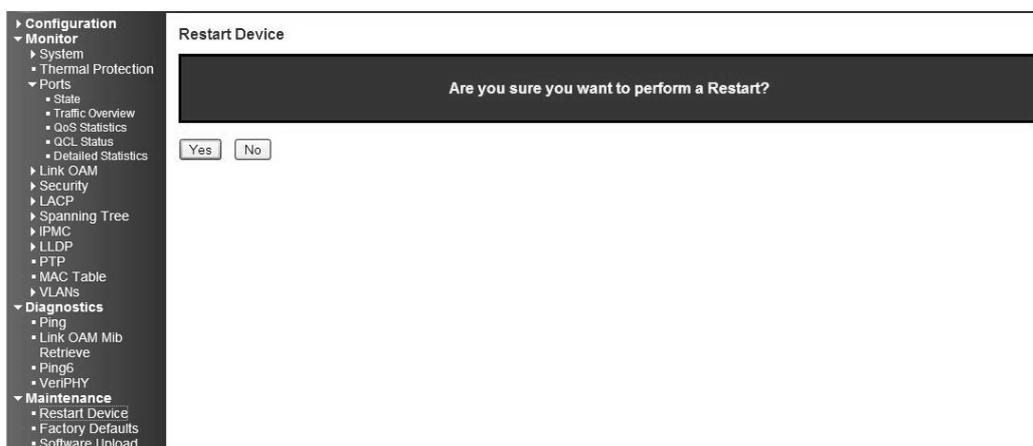
Message: Factory Defaults - Are you sure you want to reset the configuration to Factory Defaults? Yes No



Meaning: Confirmation message displayed when you select **Maintenance > Factory Defaults**.

Recovery: None; confirm that you want to reset the LIB-44xx to its factory default settings (click **Yes**), or clear the message and continue using the current configuration (click **No**).

Message: Restart Device - Are you sure you want to perform a Restart? Yes No



Meaning: Confirmation message; you selected **Maintenance > Restart Device**.

Recovery: None; confirm that you want to restart (soft boot) the LIB-44xx (click **Yes**), or clear the message and continue operation (click **No**).

Message: Invalid Firmware Image - The uploaded firmware image is invalid. Please use a correct firmware image.



Meaning: At **Maintenance > Software Upload** you entered or selected an unacceptable file (image).

Recovery:

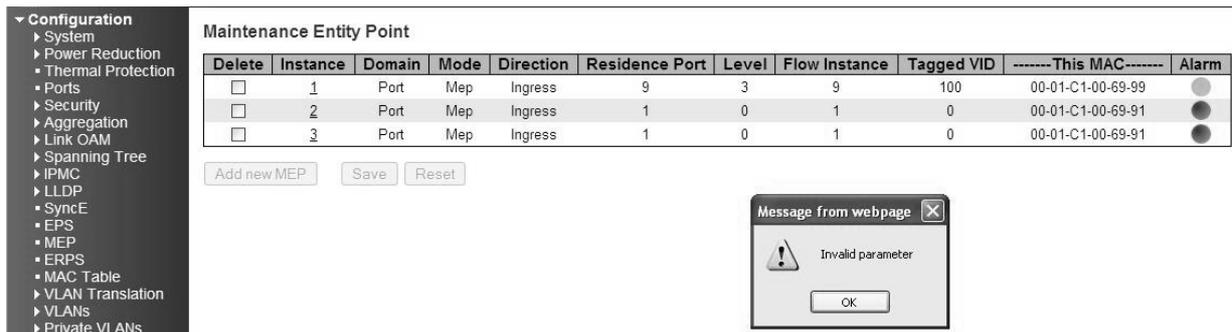
Click the browser’s **Back** key to return to the main menu.

Enter or Browse to and select an acceptable file (image) from a valid location (e.g., C:\ TFTP-Root).

Click the **Upload** button. See the “Software Upload” section for more information.

If the problem persists, contact Tech Support. See the “Service” section on page 512.

Message: Invalid parameter



Meaning:

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Change the Instance Number or Residence Port entry.
3. See “MEP Instance Configuration” on page 161.
4. If the problem persists, contact Tech Support. See the “Service” section on page 512.

Message: Do you want to log out the web site?



Meaning: You clicked the Logout () button.

Recovery: Confirmation message only; click the webpage **OK** button to continue to log out of the LIB-44xx web interface session, otherwise, click the **Cancel** button to continue working in the current LIB-44xx web session.

If you click the **OK** button to log out, the session terminates (ends) and the Connect to screen displays:



You can log in again at any time.

Message: W_port and E_port cannot be the same

The screenshot shows the 'Ethernet Ring Protection Switching' configuration page. A table lists configuration parameters for two protection groups. The first group has ERPS ID 1, Port 0 8, Port 1 7, Port 0 SF MEP 1, Port 1 SF MEP 2, Port 0 APS MEP 1, and Port 1 APS MEP 2. The second group has ERPS ID 0, Port 0 1, Port 1 1, Port 0 SF MEP 1, Port 1 SF MEP 1, Port 0 APS MEP 1, and Port 1 APS MEP 1. A message box is displayed with the text: 'W_port and E_port can not be same'.

Meaning:

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Change the **Port 0** or the **Port 1** field so that they have different entries.
3. Click the **Save** button.
4. If the problem persists, contact Tech Support. See the “Service” section on page 512.

Message: West RAPS MEP and East RAPS MEP cannot be the same

The screenshot shows the 'Ethernet Ring Protection Switching' configuration page. A table lists configuration parameters for two protection groups. The first group has ERPS ID 1, Port 0 8, Port 1 7, Port 0 SF MEP 1, Port 1 SF MEP 2, Port 0 APS MEP 1, and Port 1 APS MEP 2. The second group has ERPS ID 0, Port 0 1, Port 1 2, Port 0 SF MEP 1, Port 1 SF MEP 1, Port 0 APS MEP 1, and Port 1 APS MEP 1. A message box is displayed with the text: 'West RAPS MEP and East RAPS MEP can not be same'.

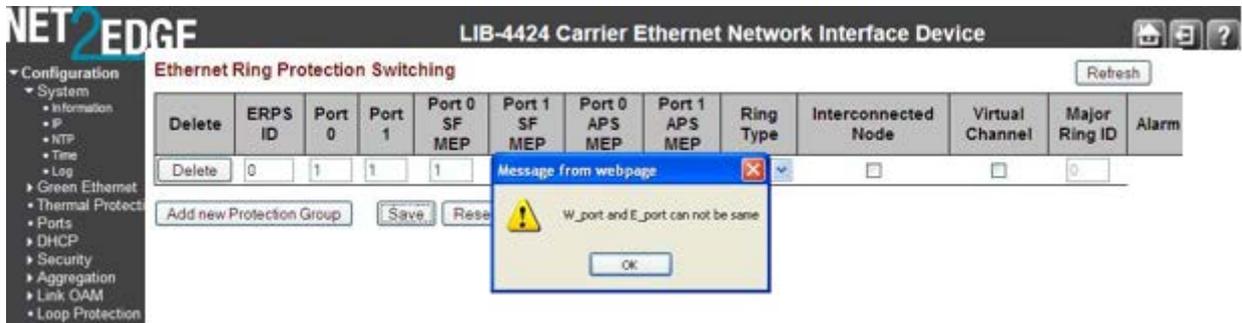
Meaning: At **Configuration > ERPS** you tried to add a new Protection Group, but the operation failed when you clicked the **Save** button.

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Change the **Port 0 SF MEP** or the **Port 1 SF MEP** field so that they have different entries.
3. Click the **Save** button.
4. If the problem persists, contact Tech Support. See the “Service” section on page 512.

Message:

West MEP and East MEP cannot be the same
W_port and E_port cannot be the same

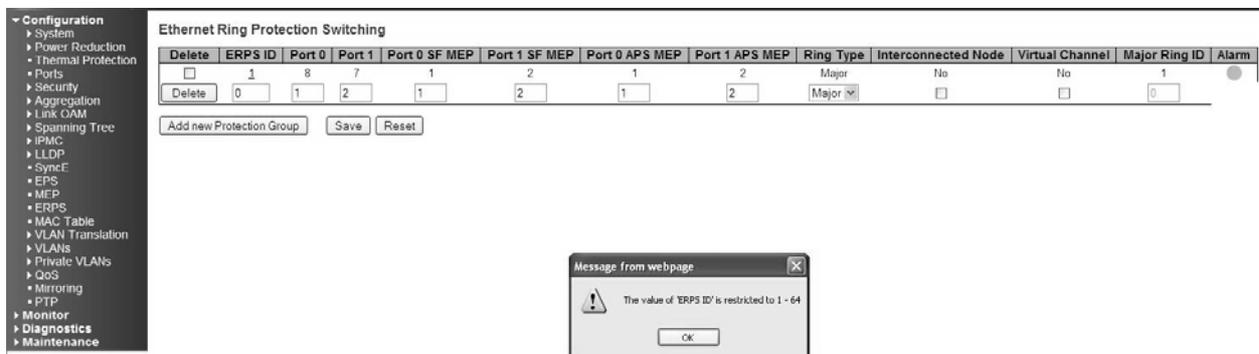


Meaning: At **Configuration > ERPS** you tried to add a new Protection Group, but the operation failed when you clicked the **Save** button.

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Change the **Port 0 APS MEP** or the **Port 1 APS MEP** field so that they have different entries.
3. Click the **Save** button.
4. If the problem persists, contact Tech Support.

Message: **The value of ERPS ID is restricted to 1-64**

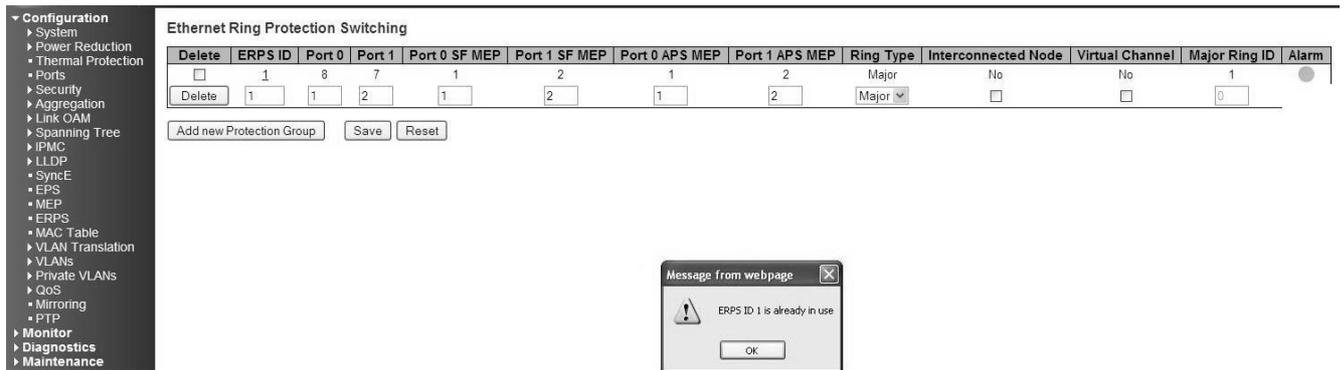


Meaning: At **Configuration > ERPS** you tried to add a new Protection Group, but the operation failed when you clicked the **Save** button.

Recovery:

1. Click the **OK** button to clear the webpage message.
2. In the **ERPDS ID** field, enter a unique ID number of 1-64.
3. Click the **Save** button.
4. If the problem persists, contact Tech Support. See the “Service” section on page 512.

Message: ERPDS ID 1 is already in use

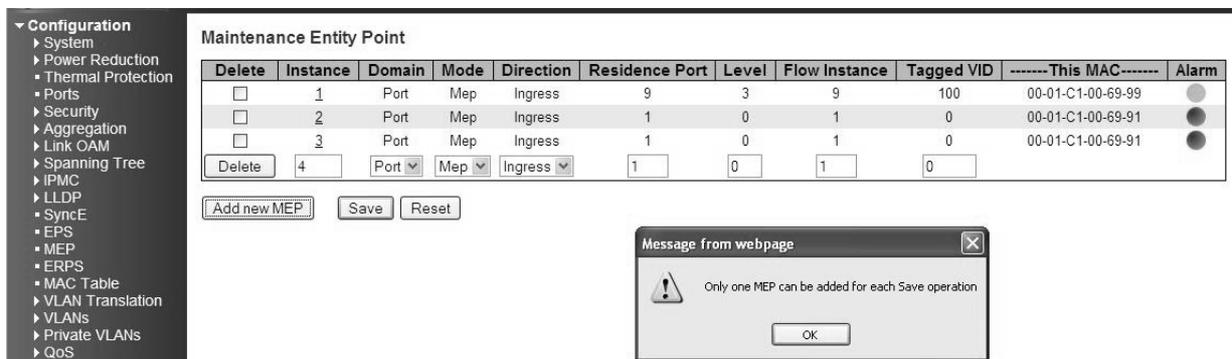


Meaning: At **Configuration > ERPS** you tried to add a new Protection Group, but the operation failed when you clicked the **Save** button.

Recovery:

1. Click the **OK** button to clear the webpage message.
2. In the **ERPDS ID** field, enter a unique ID number of 1-64.
3. Click the **Save** button.
4. If the problem persists, contact Tech Support.

Message: Only one MEP can be added for each Save operation



Meaning: At **Configuration > MEP** you clicked **Add New MEP** before you saved the MEP you were already adding.

Recovery:

1. Click **OK** to clear the webpage message.
2. Click the **Save** button, and make the required MEP config selections.
3. Click the **Add a new MEP** button to configure the next new MEP.
4. If the problem persists, contact Tech Support. See the “Service” section on page 512.

Message: Invalid parameter

Example:

- ▼ Configuration
- ▶ System
- ▶ Power Reduction
- ▶ Thermal Protection
- ▶ Ports
- ▶ Security
- ▶ Aggregation
- ▶ Link OAM
- ▶ Spanning Tree
- ▶ IPMC
- ▶ LLDP
- ▶ SyncE
- ▶ EPS
- ▶ MEP
- ▶ ERPS
- ▶ MAC Table
- ▶ VLAN Translation
- ▶ VLANs
- ▶ Private VLANs
- ▶ QoS
- ▶ Mirroring

Maintenance Entity Point

Delete	Instance	Domain	Mode	Direction	Residence Port	Level	Flow Instance	Tagged VID	-----This MAC-----	Alarm
<input type="checkbox"/>	1	Port	Mep	Ingress	9	3	9	100	00-01-C1-00-69-99	
<input type="checkbox"/>	2	Port	Mep	Ingress	1	0	1	0	00-01-C1-00-69-91	
<input type="checkbox"/>	3	Port	Mep	Ingress	1	0	1	0	00-01-C1-00-69-91	
<input type="checkbox"/>	4	Port	Mep	Ingress	1	0	1	0	00-01-C1-00-69-91	
<input type="checkbox"/>	5	Port	Mep	Ingress	1	0	1	0	00-01-C1-00-69-91	

Message from webpage ✕

Invalid parameter

Meaning: At **Configuration > MEP > Add New MEP** you selected an unsupported feature (e.g., Mode = MIP).

Recovery:

1. Click **OK** to clear the webpage message.
2. Click the **Add a new MEP** button, make another (valid / supported) selection, and then click **Save**.
3. If the problem persists, contact Tech Support. See the “Service” section on page 512.

Message: Group ID 3 and VID: 4 combination is already in use. Please update existing entry instead of adding it again!

Example:

- ▼ Configuration
- ▶ System
- ▶ Power Reduction
- ▶ Thermal Protection
- ▶ Ports
- ▶ Security
- ▶ Aggregation
- ▶ Link OAM
- ▶ Spanning Tree
- ▶ IPMC
- ▶ LLDP
- ▶ SyncE
- ▶ EPS
- ▶ MEP
- ▶ ERPS
- ▶ MAC Table
- ▼ VLAN Translation
 - ▶ Port to Group Mapping
 - ▶ VID Translation Mapping
- ▶ VLANs

VLAN Translation Table

Delete	Group ID	VLAN ID	Translated to VID
Delete	1	2	3
Delete	2	3	4
Delete	3	4	5
Delete	3	4	

Auto-refresh

Message from webpage ✕

Group ID 3 and VID: 4 combination is already in use. Please update existing entry instead of adding it again!

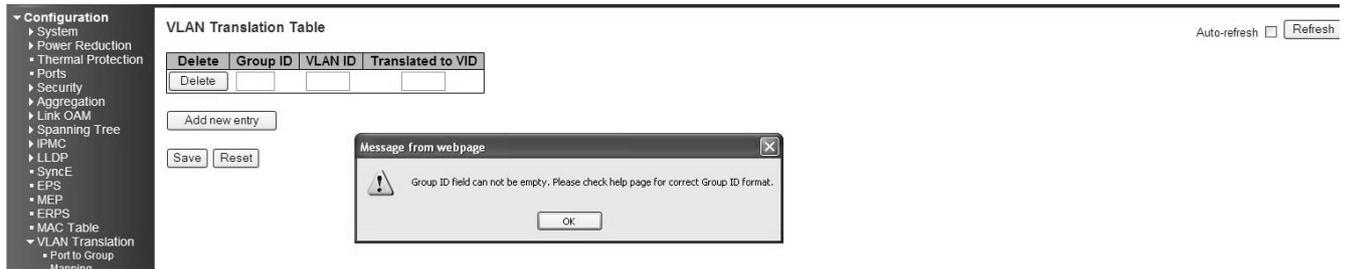
Meaning: At **VLAN Translation > VID Translation Mapping**, you tried to enter duplicate information.

Recovery:

1. Make sure the **Group ID**, **VLAN ID**, and **Translated to VID** entries are unique (do not already exist in the VLAN Translation Table).
2. Continue the operation.
3. If the problem persists, contact Tech Support. See the “Service” section on page 512.

Message: Group ID field cannot be empty. Please check help page for correct Group ID format.

Example:



Meaning:

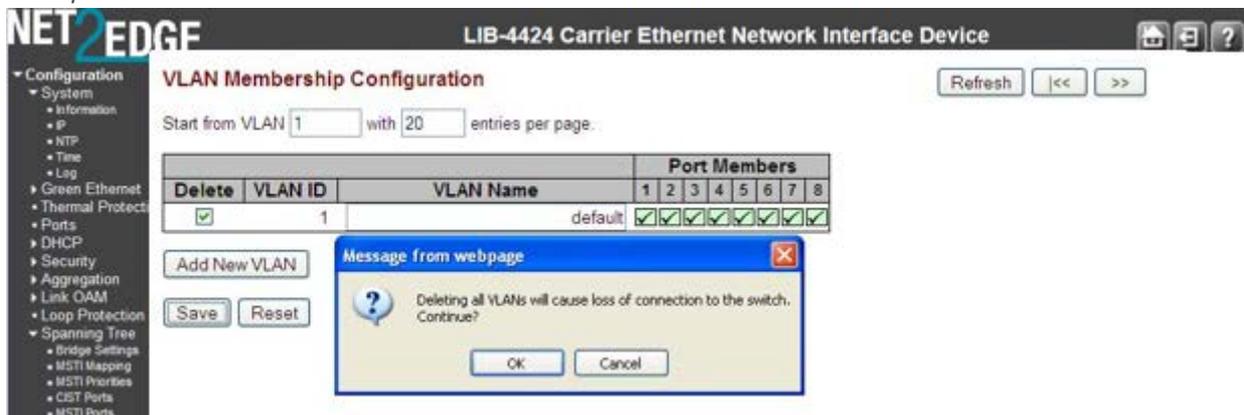
At **VLAN Translation > VID Translation Mapping**, you tried to Add a new entry, but did not enter a Group ID.

Recovery:

1. Make sure the **Group ID, VLAN ID, and Translated to VID** entries are made in the VLAN Translation Table).
2. Continue the operation.
3. If the problem persists, contact Tech Support. See the “Service” section on page 512.

Message: Deleting all VLANs will cause loss of connection to the switch. Continue?

Example:



Meaning: At **Configuration > VLANs > VLAN Membership**, you tried to delete VLAN ID 1 from the ‘VLAN Membership’ table. This is the default VLAN; if you delete it the LIB-44xx connection will drop.

Meaning: You tried to delete a non-existent VLAN translation entry from a group.

Recovery:

1. Make sure this is the action you want. If so, click the **OK** button; if not, click the **Cancel** button to clear the webpage message. Continue operation.
2. Re-enter the command with a different (existing) group.
3. Add port member to the group and re-try the operation.
4. Make sure you are not trying to delete VLAN 1. Deleting VLAN 1 causes issues with forwarding.
5. To make sure no ports belong to VLAN 1, add VLAN 1 with all ports in the forbidden state.

6. See the **Delete VLAN Translation Group Entry** command for CLI information.
7. If a problem persists, contact Tech Support. See the “Service” section on page 512.

Message: The value of ‘Queue Shaper Rate’ is restricted to 100 - 1000000 kbps. Select the ‘Mbps’ unit for coarser granularity.

Example:



Meaning: At **Configuration > QoS > Port Scheduler**, you selected too high or low a number for the unit of measure. You selected too fine of a granularity (measurement)

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Select another unit of measure (‘kbps’ or ‘mbps’) and continue operation.
3. If a problem persists, contact Tech Support. See the “Service” section on page 512.

Message: Do you want to log out the web site?

Example:



Meaning: FireFox message that displays when you click the web interface **Logout** button.

Recovery:

1. Click the **OK** button to clear the webpage message. The login prompt displays again.
2. Continue operation - if desired, log back in to the LIB-44xx web interface.
3. If a problem persists, contact Tech Support.

**Message: The value of Port Error Recovery Timeout is 30-84600
The value of Port Error Recovery Timeout cannot be empty**

Example:



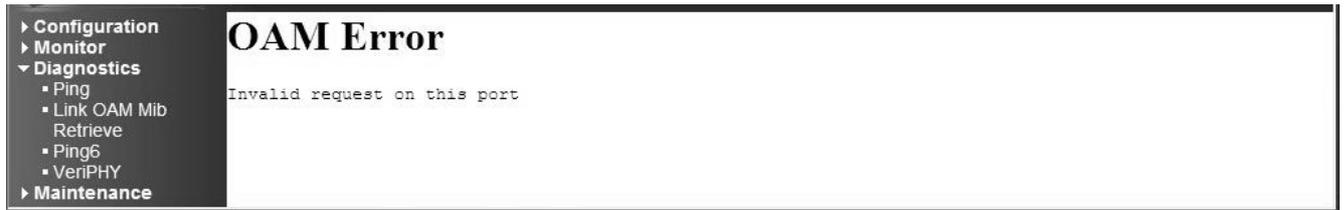
Meaning: At **Configuration > Spanning Tree > Bridge Settings** you entered an invalid value in the “Port Error Recovery Timeout” field.

Recovery:

1. Enter a value in the range of 30-84,600 seconds (½ minute - 1243.3 minutes or 20.72 hours).
2. Continue operation.
3. If a problem persists, contact Tech Support.

Message: **OAM Error - Invalid request on this port**

Example:



Meaning:

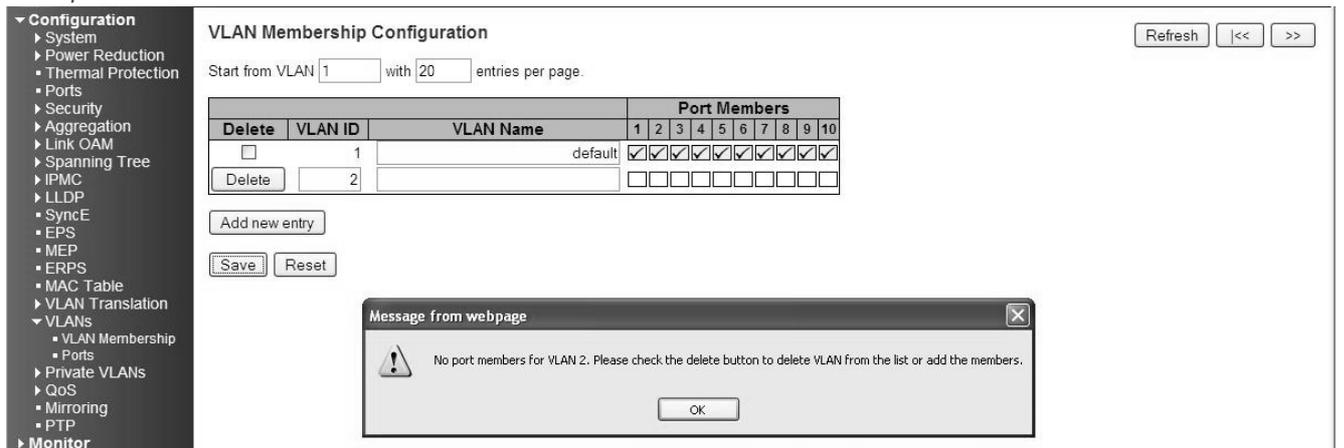
At **Diagnostics > Link OAM Mib Retrieve** you selected a port that is not ...

Recovery:

1. Click the browser’s back button to clear the message, verify your selection. .
2. Make sure you have selected either the “Local” or “Peer” radio button.
3. Verify the port number selected.
4. If a problem persists, contact Tech Support.

Message: **No port members for VLAN x. Please check the delete button to delete the VLAN from the list or add the members.**

Example:



Meaning: At **Configuration > VLANs > VLAN Membership** you tried to add a VLAN to the VLAN membership configuration, but you did not check any of the Port Members checkboxes.

Note: VLAN 1 cannot be deleted.

Recovery:

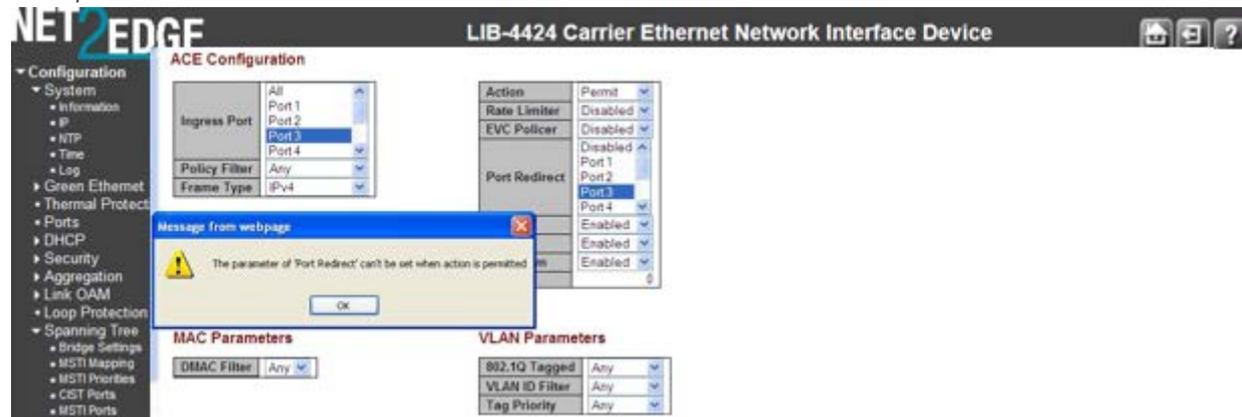
1. Click the **OK** button to clear the webpage message.
2. Either a) click the **Delete** button to delete the VLAN from the list, or b) click one or more Port Members checkboxes in the table and click the **Save** button to add the new entry.
3. If a problem persists, contact Tech Support. See the “Service” section on page [512](#).

Message: The parameter of ‘port_copy’ can’t be set when action is permitted

Meaning: At **Configuration > Security > Network > ACL > Ports** you set the “Action” parameter to “Permit”, which does not allow the “Port Copy” parameter to be set to a Port (**Port 3** on the screen above).

Recovery:

1. Click the **OK** button to close the webpage message.
2. Either a) change the “Action” parameter selection, or b) change the “Port Copy” parameter selection to “Disabled”.
3. Click the **Save** button when done.
4. If a problem persists, contact Tech Support. See the “Service” section on page [512](#).

Message: The parameter of ‘Port Redirect’ can’t be set when action is permitted*Example:*

Meaning: At **Configuration > Security > Network > ACL > Access Control List** you set the “Action” parameter to “Permit”, which does not allow the “Port Redirect” parameter to be set to a Port (**Port 3** on the screen above).

Recovery:

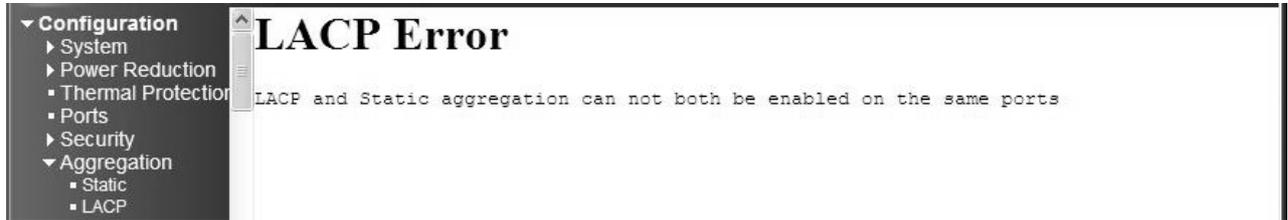
1. Click the **OK** button to close the webpage message.
2. Either a) change the “Action” parameter selection, or b) change the “Port Redirect” parameter selection to “Disabled”.
3. Click the **Save** button when done.
4. If a problem persists, contact Tech Support. See the “Service” section on page [512](#).

Message: Aggregation Error - Port joining aggregation must be in the same speed and in full duplex

Group 1 member counts error!! Local aggregation must include 2-16 ports.

LACP Error - LACP and Static aggregation can not both be enabled on the same ports

Example:



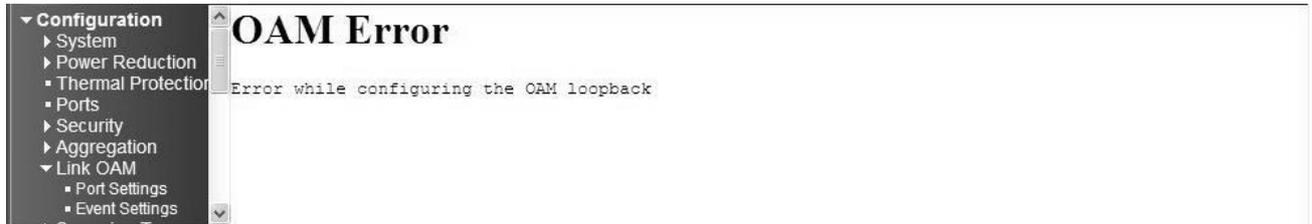
Meaning: You configured a port for both LACP and Static aggregation, which is not supported. For example, at **Configuration > Aggregation > Static**. and at **Configuration > Aggregation > LACP** you configured a port for both LACP and Static aggregation, which is not supported.

Recovery:

1. Click the browser 'Back' button to clear the message.
2. Make sure each port has one configuration (either LACP or Static aggregation) enabled.
3. Verify the aggregation path configuration, click Save, and continue operation. See the "[Aggregation Configuration](#)" section on page 63.
4. If a problem persists, contact Tech Support. See the "Service" section on page 512.

Message: **OAM Error - Error while configuring the OAM loopback**

Example:



Meaning: At **Configuration > Link OAM > Port Settings** you entered an invalid parameter.

Recovery:

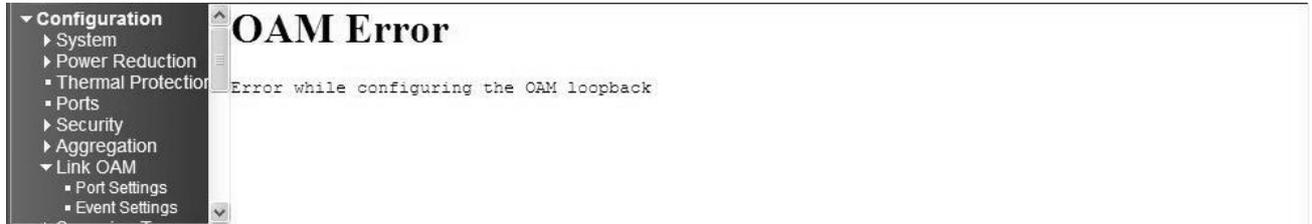
1. Click the browser 'Back' button to clear the message.
2. Make sure the Loopback Operation checkbox is unchecked.
3. Verify the 'Link OAM Port Configuration' table selections, click **Save**, and continue operation. See the "

[Link OAM \(LOAM\) Configuration](#)” section on page 112.

4. If a problem persists, contact Tech Support. See the “[Service](#)” section on page 512.

Message: **OAM Error - Error While Configuring Link Events**

Example:



Meaning: At **Configuration > Link OAM > Event Settings** you entered an invalid parameter.

Recovery:

1. Click the browser ‘Back’ button to clear the message.
2. Verify the ‘Link Event Configuration for Port x’ selections, click ‘**Save**’, and continue operation.
3. If a problem persists, contact Tech Support. See the “[Service](#)” section on page 512.

Message: **HW Install - cannot start with this user account.** Make sure this user account is a member of the Administrators group on this computer.

You do not have sufficient security privileges to install hardware on this computer. Please contact your site administrator, or logout and log in again as an administrator and try again.

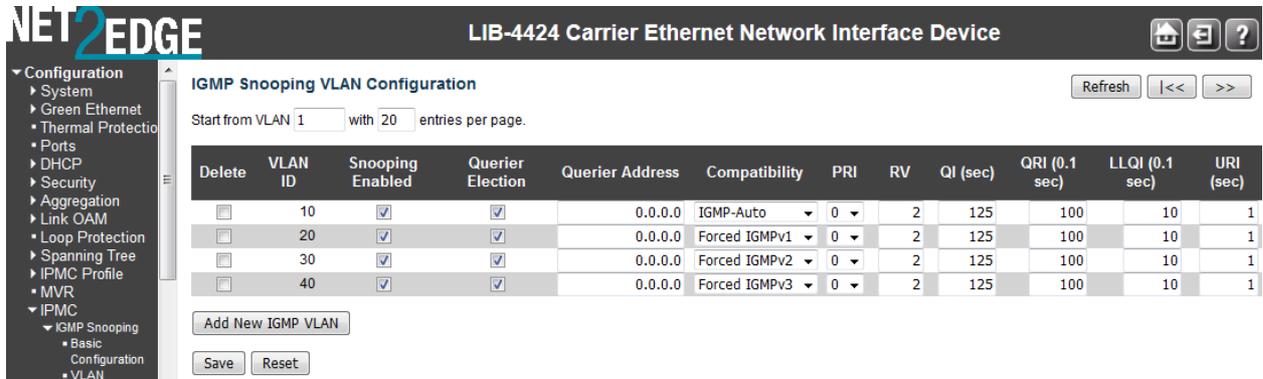


Meaning: A privilege level issue exists.

Recovery:

1. Contact your system administrator, or change your user privilege to Admin.
2. Click the **OK** button to clear the message.
3. Continue operation.
4. If a problem persists, contact Tech Support. See the “[Service](#)” section on page 512.

Message: The value of ‘Group IP Address’ must be a valid IP address in dotted decimal notation (x.y.z.w).



Meaning: At **Configuration > IPMC > IGMP Snooping > Port Group Filtering**, you entered an invalid IP address.

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Make sure the IP Address you enter is in dotted decimal notation (x.y.z.w), where: **x** is a decimal number from 224 to 239, and **y**, **z**, and **w** are decimal numbers from 0 to 255.
3. Continue operation. See the “[Port Group Filtering](#)” section on page 100.
4. If a problem persists, contact Tech Support. See the “[Service](#)” section on page 512.

Message: Using IPv6 unicast address is not allowed here.



Meaning: At **Configuration > IPMC > MLD Snooping > Port Group Filtering** you entered an IPv6 unicast address in the Filtering Groups field.

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Make sure the IP Address you enter is an IPv6 multicast address.
3. Continue operation. See the “[Port Group Filtering](#)” section on page 100.
4. If a problem persists, contact Tech Support. See the “[Service](#)” section on page 512.

Message: The value of ‘Group Address’ must be a valid IPv6 address in 128-bit records represented as eight fields of up to four hexadecimal digits with a colon separates each field (:).

Meaning: At **Configuration > IPMC > MLD Snooping > Port Group Filtering** you entered an invalid IP address in the Filtering Groups table / field.

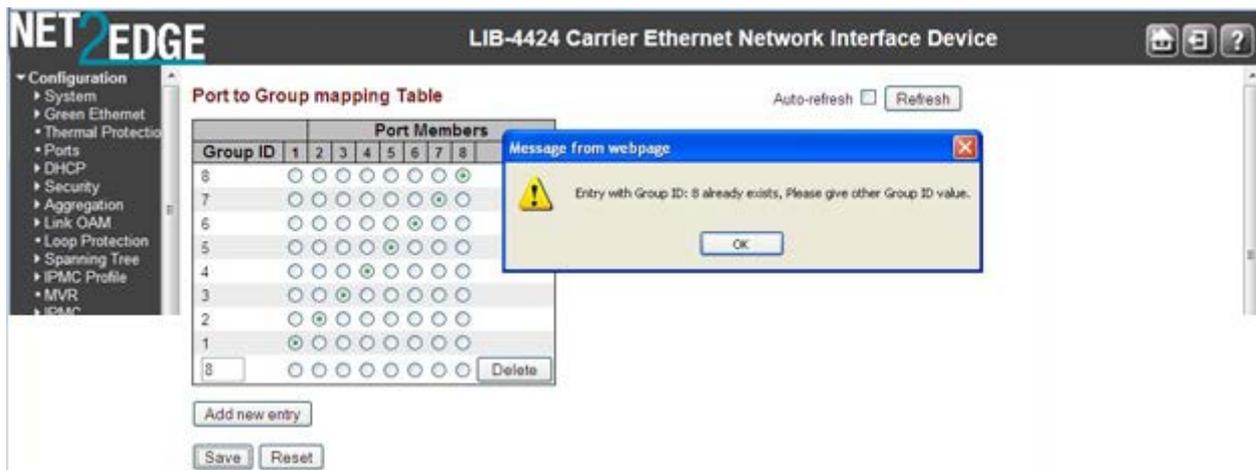
Meaning: At **Configuration > IPMC > MLD Snooping > Port Group Filtering** you entered an IPv6 unicast address in the Filtering Groups field.

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Make sure the IP Address you enter is a valid IPv6 address in 128-bit records represented as eight fields of up to four hexadecimal digits with a colon separates each field (:).
3. Continue operation. See the “[Port Group Filtering](#)” section on page 100.
4. If a problem persists, contact Tech Support. See the “[Service](#)” section on page 512.

Message:

Entry with Group ID: x already exists, Please give other Group ID value.
Invalid Group ID value: xx. Group ID must be an integer between 1 to 8.



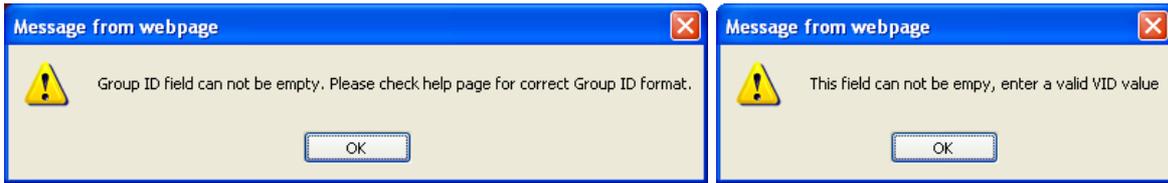
Meaning: At **Configuration > VLAN Translation** you entered a Group ID number outside of the valid range of 1-8, or you entered an existing Group ID number.

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Enter a valid, unique Group ID number and click the **Save** button. See “[Port to Group Mapping](#)” on page 131.
3. Verify that the Add new entry was successful, and continue operation.
4. If a problem persists, contact Tech Support. See the “[Service](#)” section on page 512.

Message: **This field can not be empty, enter a valid VID value.**

Example: Group ID field cannot be empty. Please check help page for correct Group ID format.

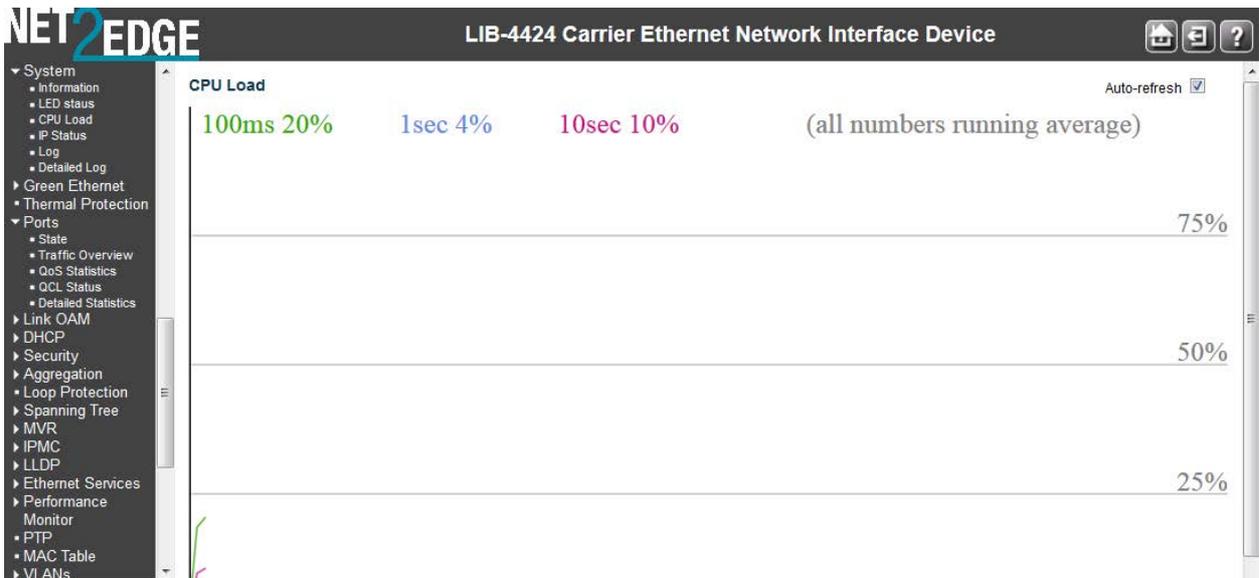


Meaning:

Recovery: At **Configuration > VLAN > VLAN Membership** you did not enter a Group ID or VLAN ID.

1. Click the **OK** button to clear the webpage message.
 2. Enter a valid, unique Group ID number or VLAN ID number and click the **Save** button.
- See “[IP Configuration](#)” on page 12.
3. Verify that the entry was successful, and continue operation.
 4. If a problem persists, contact Tech Support. See the “[Service](#)” section on page 512.

Message: **Collecting initial data, please wait ..**



Meaning: Displayed in Google Chrome or Mozilla Firefox at the **Monitor > System > CPU Load** menu path. This page displays the CPU load, using an SVG graph. In order to display the SVG graph, your browser must support the SVG format.

Recovery:

1. See the SVG Wiki at http://wiki.svg.org/index.php/Viewer_Implementations for more information on browser support.
2. At the [SVG Wiki](#), check for native implementations and/or download available Browser Plug-Ins.
3. Retry the operation. If necessary, use Internet Explorer as your web browser.
4. If a problem persists, contact Tech Support.

Message: **The value of Path Cost cannot be empty
The value of 'Path Cost' is restricted to 1 – 20000000**

STP CIST Port Configuration

CIST Aggregated Port Configuration

Port	STP Enabled	Path Cost	Priority	Admin Edge	Auto Edge	Restricted Role	TCN	BPDU Guard	Point-to-point
-	<input type="checkbox"/>	Auto	128	Non-Edge	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Forced True

CIST Normal Port Configuration

Port	STP Enabled	Path Cost	Priority	Admin Edge	Auto Edge	Restricted Role	TCN	BPDU Guard	Point-to-point
*	<input type="checkbox"/>	<>	<>	<>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<>
1	<input type="checkbox"/>	Auto	128	Non-Edge	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Auto
2	<input type="checkbox"/>	Auto	128	Non-Edge	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Auto
3	<input type="checkbox"/>	Auto	128	Non-Edge	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Auto
4	<input type="checkbox"/>	Auto	128	Non-Edge	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Auto
5	<input type="checkbox"/>	Auto	128	Non-Edge	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Auto
6	<input type="checkbox"/>	Auto	128	Non-Edge	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Auto

Message from webpage

The value of 'Path Cost' cannot be empty.

OK

Meaning: At **Configuration > Spanning Tree > CIST Ports** you selected “**Specific**” in the **Path Cost** dropdown, but did not enter a value in the **Path Cost** entry field.

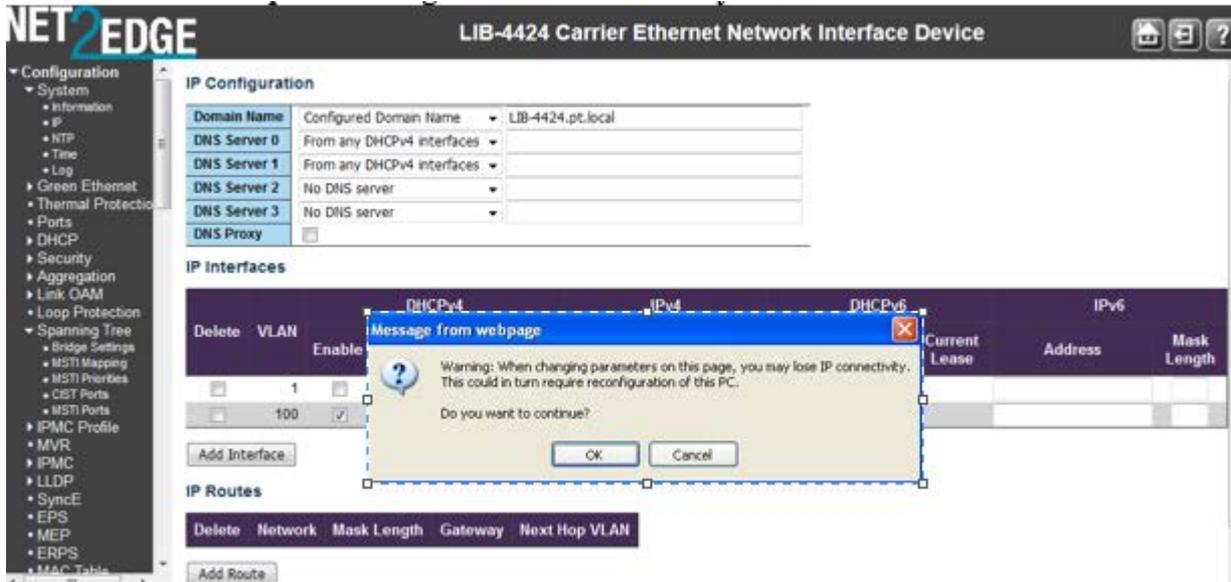
Recovery:

1. Click **OK** to clear the webpage message.
2. Either enter a valid **Path Cost** (1 - 200,000,000) in the entry field, or select **Auto** at the dropdown.

See “[CIST Ports](#)” on page 124.

3. If a problem persists, contact Tech Support. See the “[Service](#)” section on page 512.

Message: Warning: When changing parameters on this page, you may lose IP connectivity. This could in turn require reconfiguration of this PC. Do you want to continue?

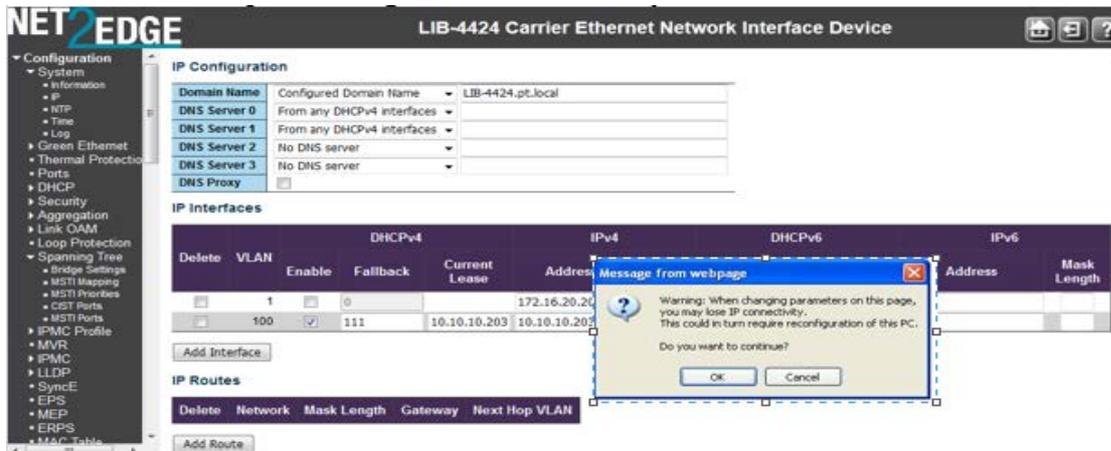


Meaning: At **Configuration > System > IPv6** you checked the **Auto Configuration Configured** checkbox and clicked **Save**.

Recovery:

1. Verify that this is the function that you want to perform and be aware of the consequences mentioned.
2. Determine whether your particular computer will require reconfiguration. For example, for Microsoft .NET Framework version 2.0 and later, IPv6 is enabled by default. For .NET Framework version 1.1 and earlier, IPv6 is disabled by default. For more information see the MSDN article at <http://msdn.microsoft.com/en-us/library/8db2058t.aspx>. Windows Server 2008 provides complete support for IPv6 and all of its features, and does not need additional installation or configuration. For Windows 7 see <http://technet.microsoft.com/en-us/magazine/2009.07.cableguy.aspx>. For Windows XP see <http://support.microsoft.com/kb/2478747>. For Windows Vista see <http://ipv6.com/articles/general/IPv6-Microsoft-Vista.htm>. For Linux / BSD, see <http://ipv6.com/articles/applications/Linux-and-BSD.htm> or your distribution documentation and/or website.
3. Click the **OK** button only if you are sure you want to continue IPv6 Auto Configuration. Otherwise click the **No** button and click **Reset**.
4. See "IPv6 Configuration" on page 15 for more information.

Message: Warning: When changing parameters on this page, you may lose IP connectivity. This could in turn require reconfiguration of this PC. Do you want to continue?

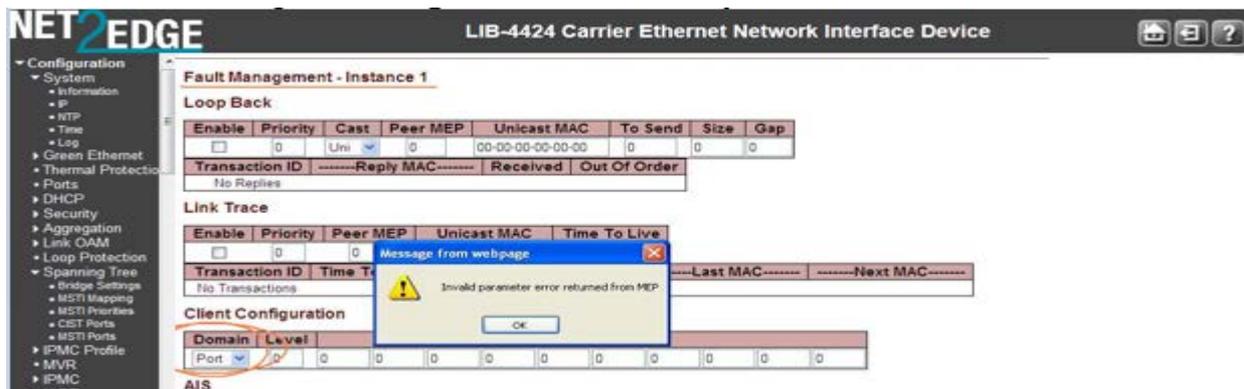


Meaning: At **Configuration > System > IP** you checked the **DHCP Client Configured** checkbox and clicked **Save**.

Recovery:

1. Verify that this is the function that you want to perform and be aware of the consequences mentioned.
2. Click the **OK** button only if you are sure you want to continue IPv6 Auto Configuration. Otherwise click the **No** button and then click **Reset**.
3. See “[IPv4 Configuration](#)” on page 12 for more information.

Message: Invalid parameter error returned from MEP



Meaning: At **Configuration > MEP > MEP Configuration > Fault Management - Instance** in the **Client Configuration** section at the **Domain** dropdown, you selected a currently unsupported client layer domain.

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Select 'EVC'. Only a 'Port' MEP is able to be a server MEP having relation to a client layer. The other selections (Esp, Evc, Mpls) are for future use.
3. See ‘[MEP Configuration](#)’ on page 156.

Message: The value of 'Gap' is restricted to 10-65535

Meaning: At **Configuration > MEP > Performance Monitor - Instance** in the **Delay Measurement** section in the **Gap** field, you entered an invalid value.

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Enter a Gap value of 10-65535.
3. See '[MEP Configuration](#)' on page 156.

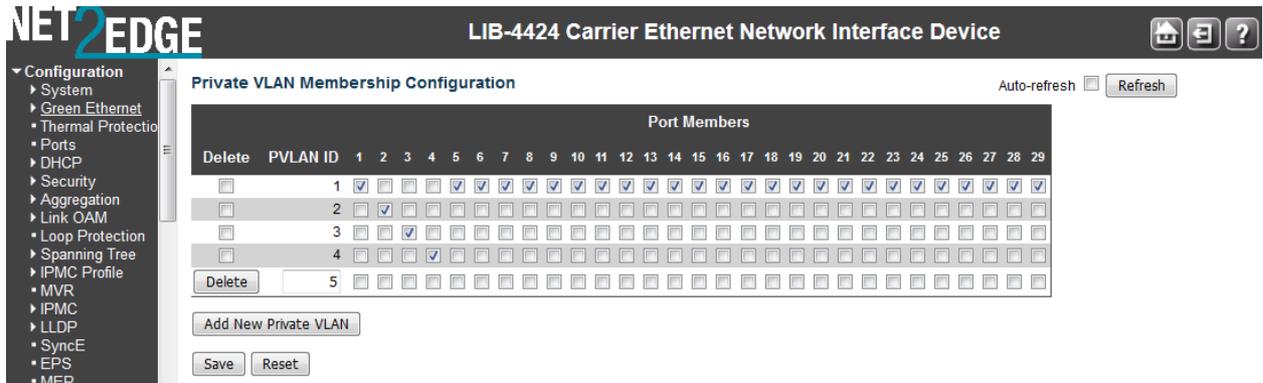
Message: The value of "Count" is restricted to 10-2000

Meaning: At **Configuration > MEP > Performance Monitor - Instance** in the **Delay Measurement** section in the **Count** field, you entered an invalid value.

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Enter a Gap value of 10-2000.
3. See '[MEP Configuration](#)' on page 156.

Message: At least one port must be selected to add an entry
At least one port must be selected. To delete entry, check the delete checkbox.

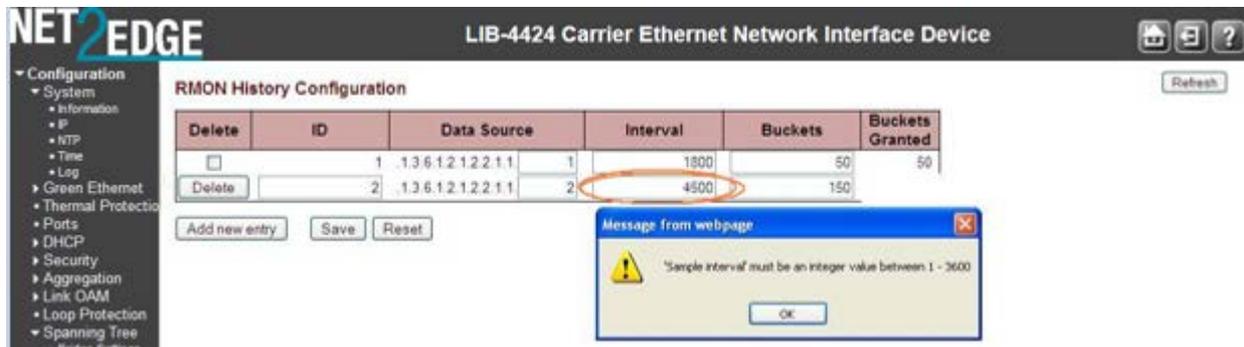


Meaning: At **Configuration > Private VLANs > PVLAN Membership** you clicked the **Save** button without first checking at least one of the **Port Membership** checkboxes.

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Check one or more of the **Port Membership** checkboxes.
3. See ‘[PVLAN Membership](#)’ on page 209.

Message: ‘Sample interval’ must be an integer value between 1- 3600

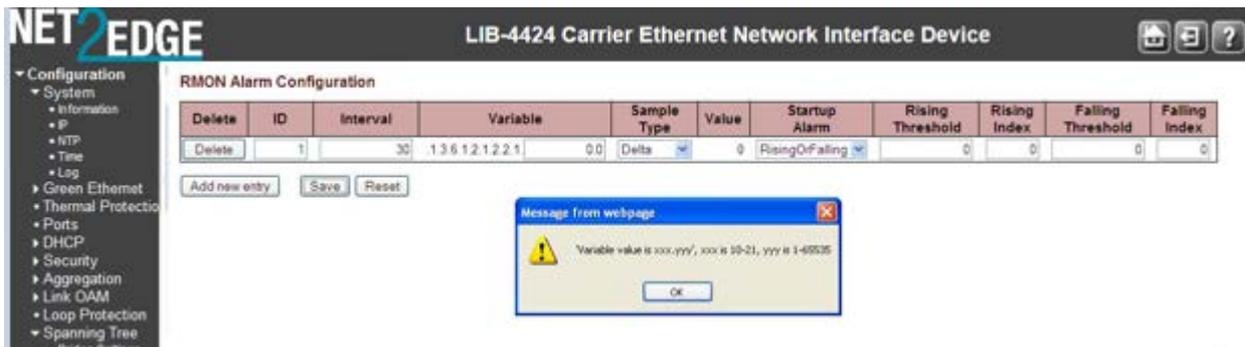


Meaning: At the **Configuration > Security > Switch > SNMP > RMON > History** menu path in the **Interval** field, you entered an invalid number.

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Enter a valid Interval in the range of 1 to 3600.
3. See “[Configuration > Security > Switch > SNMP > RMON](#)” on page 83.

Message: 'Variable value' is xxx.yyy', xxx is 10-21, yyy is 1-65535



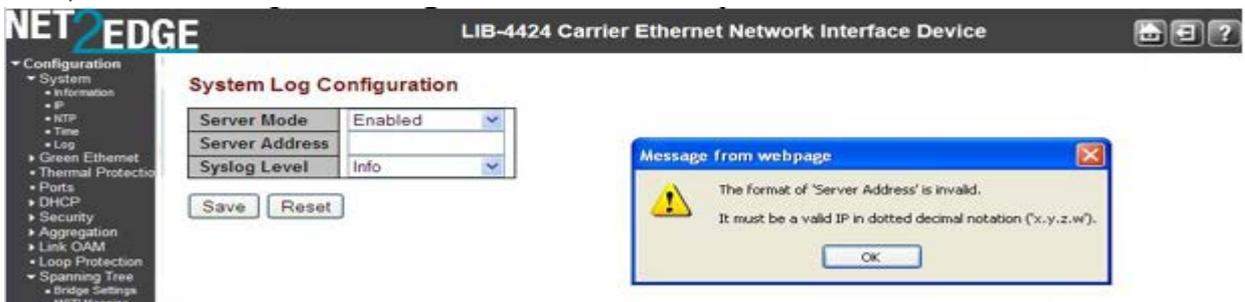
Meaning: At the **Configuration > Security > Switch > SNMP > RMON > Alarm** menu path in the **Variable** field, you entered an invalid number.

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Enter a valid Variable value.
3. See “[Configuration > Security > Switch > SNMP > RMON](#)” on page 83.

Message: The format of 'Server Address' is invalid. It must be a valid IP in dotted decimal notation 'x.y.z.w').

Example



Meaning:

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Re-enter a valid IP address in dotted decimal notation (e.g., 192.168.1.30).
3. See “[NTP Configuration](#)” on page 50.

Message: **This browser doesn't support dynamic tables.**

Meaning: The browser you are using is not supported.

Recovery: Use a browser that the LIB-44xx supports. See “[Web Browser Support](#)” on page 32.

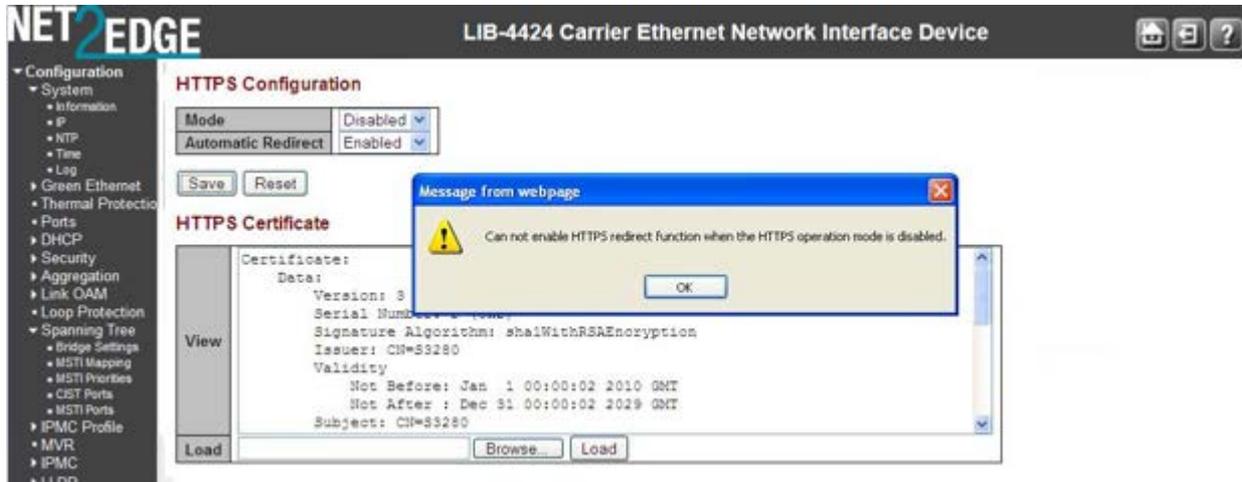
Message: **Port %lu does not support this mode\n**

Meaning: The DMI function is not supported on this port.

Recovery:

1. Switch to another port that supports DMI.
2. Use another LIB-44xx function.
3. See “[DMI Configuration](#)” on page 50.

Message: Cannot enable HTTPS redirect function when the HTTPS operation mode is disabled.

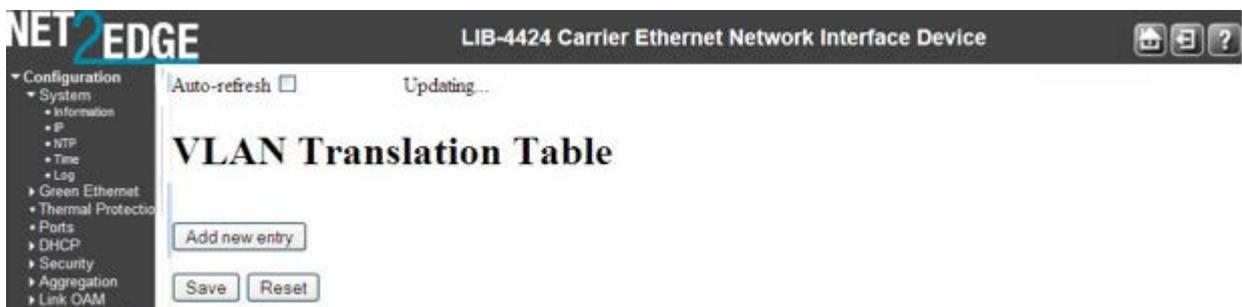


Meaning: At **Configuration > Security > Switch > HTTPS** you selected an HTTPS configuration of Mode = Disabled and Automatic Redirect = Enabled, which is not supported.

Recovery:

1. Click the **OK** button to clear the webpage message.
2. At **Configuration > Security > Switch > HTTPS** select an HTTPS configuration of either:
 - Mode = Disabled and Automatic Redirect = Disabled**, or
 - Mode = Enabled and Automatic Redirect = Disabled**, or
 - Mode = Enabled and Automatic Redirect = Enabled**.
3. See “[HTTPS Configuration](#)” on page 200 for more information.

Message: VLAN Translation Table -- Updating...



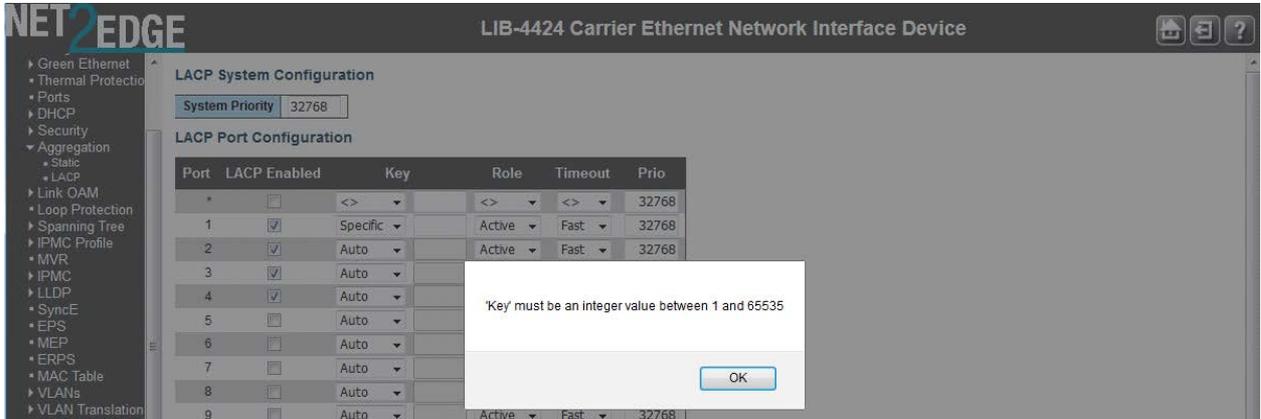
Meaning: At **Configuration > VLAN Translation > VID Translation Mapping** you tried to add a new VLAN translation table entry, but the attempt failed.

Recovery:

1. Press the **Reset** button.
2. Click the browser Back button.

3. At the CLI, enter the command **config default keep_ip** and press **Enter**.
4. Restart the LIB-44xx web interface.

Message: The value of 'Key cannot be empty



Meaning: At **Configuration > Aggregation > LACP** you clicked the **Save** button without first entering a **Key** entry field entry.

Recovery:

1. Click the **OK** button to close the message dialog box.
2. Enter a **Key** entry field entry.
3. Click the **Save** button when done. See ‘[LACP \(Link Aggregation Control Protocol\)](#)’ on page 109.

Message: Switch does not respond.



Meaning: You tried to enable both ACL policer and EVC policer functions at the same time.

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Log back in to the LIB-44xx web interface if your web browser can no longer display the webpage.
3. Retry the operation.
4. Check the CLI for the message “*Error: ACL policer and EVC policer can not both be enabled*”.
5. At the CLI, press the **Enter** key to re-display the CLI main (startup) screen.
6. Enter the CLI command “**config default keep_config**”.
7. Log in again via the LIB-44xx web interface.

Message: HTTPS Certificate Load Error. SSL Certificate PEM file size too big

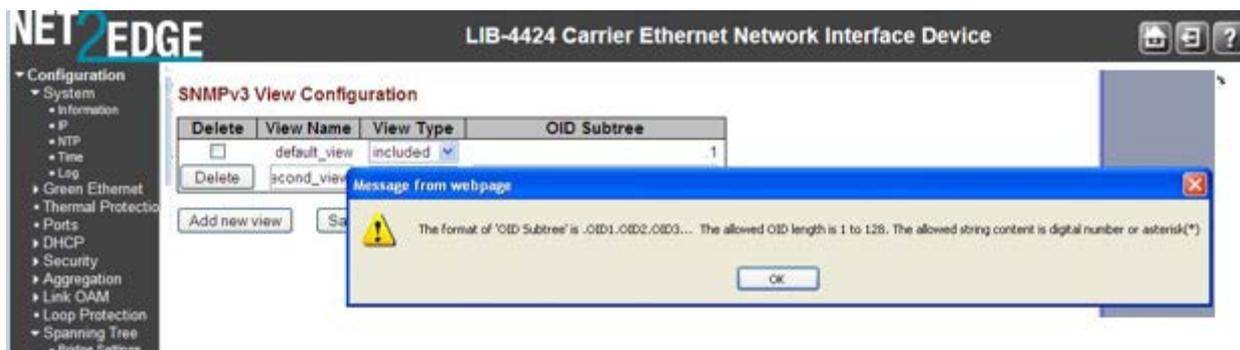


Meaning: At **Configuration > Security > Switch > HTTPS** you tried to download a HTTPS / SSL certificate file that exceeded the file size limit.

Recovery:

1. Click the browser Back button.
2. Check the SSL certificate file in terms of size and content.
3. Retry the operation. See “[HTTPS Configuration](#)” on page 61.

Message: The format of OID Subtree is .OID1.OID2.OID3. The allowed length is 1 to 128. The allowed string content is digital number or asterisk (*).

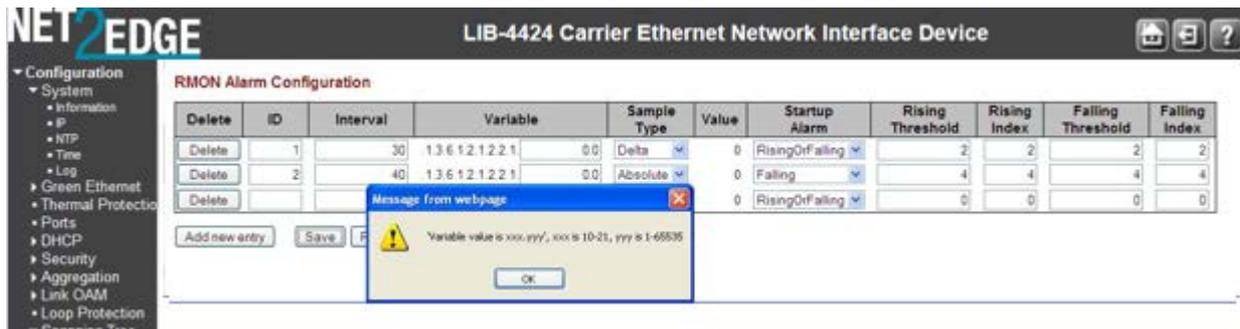


Meaning: At **Configuration > Security > Switch > SNMP > Views** you entered an invalid OID Subtree value.

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Enter an OID Subtree value in the correct format, length and content.
3. See “[SNMP Configuration](#)” on page 74.

Message: 'Variable value is xxx.yyy', xxx is 10-21, yyy is 1-65535



Meaning: At **Config > System > Security > RMON > Alarm** you entered an invalid value.

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Enter a Variable value in the correct format, length and/or content.
3. See “[SNMP Configuration](#)” on page 74.

Messages: The **HTTPS** function enabled on this device. Redirect for using **HTTPS ... Content was blocked because it was not signed by a valid security certificate.** To help protect your security, Internet Explorer has blocked this website with security certificate errors. [Click here for option.](#)



Meaning: Web browser certificate message.

Recovery:

- Click the browser’s Back button to clear the error message.
- See “Certificate Errors” in IE Help.
- Click the Information Bar (“To help protect ...”) and then click “**Display Blocked Content**”.
- At the message “There is a problem with this website’s security certificate.”, click “Continue to this website (not recommended).”
- Log back in to the LIB-44xx system.
- Click on ‘Certificate Error’.
- Click on “View Certificates”.
- Click on ‘Install Certificate...’.
- At the “Welcome to Certificate Import Wizard” - “Welcome” screen, click **Next**.
- At the Certificate Store” dialog box, click **Next**.
- At the “Completing the Certificate Import Wizard” dialog box, click the “**Finish**” button.
- The “Security Warning” dialog box displays.



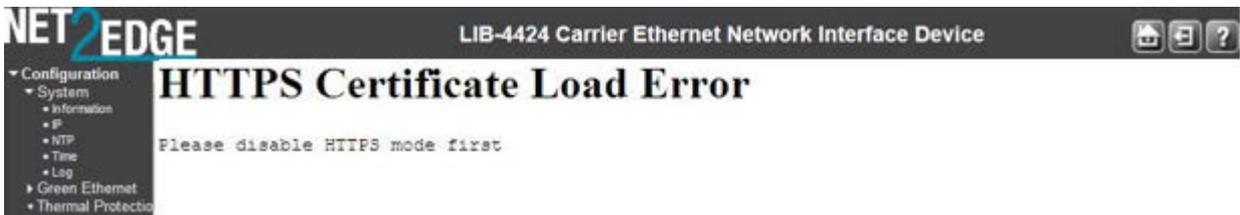
Click the **Yes** button.

When the “*Import was successful*” message displays, click the **OK** button.

At the Certificate dialog box / General tab, click the **OK** button.

Continue the operation; see “[HTTPS Configuration](#)”.

Message: HTTPS Certificate Load Error - Please disable HTTPS mode first



Meaning: At **Configuration > Security > Switch > HTTPS** you tried to load an HTTPS certificate with HTTPS mode enabled.

Recovery:

Click the browser Back button to clear the error message.

At **Configuration > Security > Switch > HTTPS** set HTTPS mode to disabled.

Continue the operation; see “[HTTPS Configuration](#)”.

Message: The Voice VLAN ID should not equal switch management VLAN ID

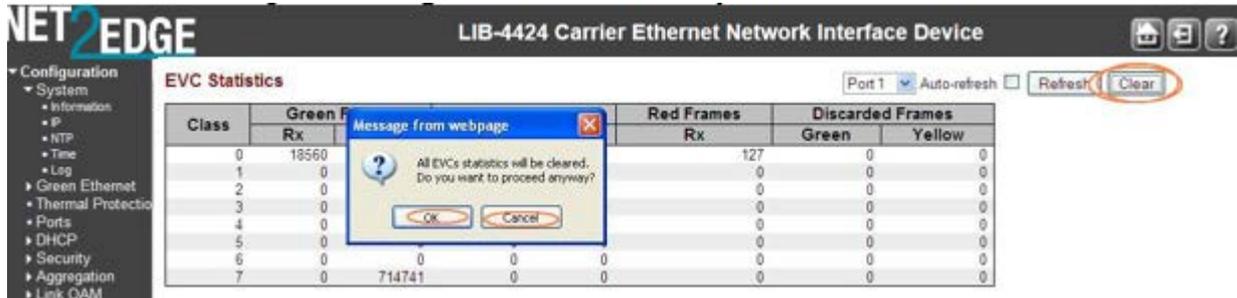


Meaning: You entered the same VLAN ID at **Configuration > MVR** and at **Configuration > Voice VLAN**.

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Change the VLAN ID at MVR or at Voice VLAN. See the related section of this manual.

Message: All EVCs statistics will be cleared. Do you want to proceed anyway?



Meaning: Confirmation message. You clicked the **Clear** button at **Monitor > Ethernet Services > EVC Statistics**.

Recovery:

1. Click the **OK** button to clear the webpage message and clear all of the stored EVC statistics, or click the **Cancel** button to leave the stored EVC statistics and continue operation.
2. See “**Monitor > Ethernet Services**” on page 390.

Message: All EVCs statistics will be removed. Do you want to proceed anyway?



Meaning: You deleted the ECE from **Configuration > Ethernet Services > ECE**, or you checked all of the UNI Ports checkboxes in the “ECE Configuration” section.

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Click the **Remove All** button. Internet connectivity is likely lost, and the LIB-44xx web interface is temporarily unavailable.
3. From the CLI, enter the command `conf default keep_ip` and press Enter.
4. Open a new web browser session and log back in to the LIB-44xx.

Message: **Frame Type: 1 and value: 0800 is already mapped to Group ID: 'a1'**



Meaning: At **Configuration > VCL > Protocol-based VLAN > Protocol to Group** you tried to **Save** a second similar configuration.

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Change the new entry's Frame Type and/or Value.
3. See [“Protocol-based VLAN”](#) on page 214.

Message: **Port Security Error** - The 802.1X Admin State must be set to Authorized for ports that are enabled for LACP

Port Security Error - The 802.1X Admin State must be set to Authorized for ports that are enabled for Spanning Tree.

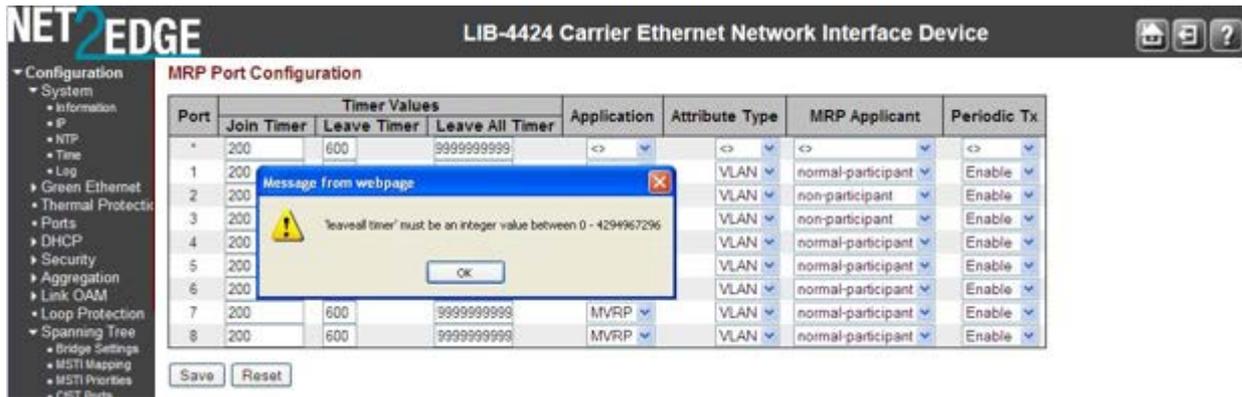


Meaning: At **Security > Network > NAS** you entered an unsupported configuration and clicked the **Save** button. In the 'Port Configuration' table in the 'Admin State' column, the parameter you selected (e.g., "Port-based 802.1X") requires a configuration change before you can perform this action.

Recovery:

1. Click the browser's **Back** button to clear the message.
2. Make the requested configuration change (e.g., disable LACP for the port, or disable Spanning Tree for the port). To configure the Spanning Tree function, see **Configuration > Spanning Tree > CIST Ports > CIST Normal Port Configuration** in the "STP Enabled" column.
3. Re-try the configuration at **Security > Network > NAS**. See [Configuration > Security > Network > NAS](#) for more information.

Message: **'leaveall timer' must be an integer value between 0 – 4294967296**



Meaning: At **Configuration > MRP** you entered an invalid value in the **Leave All Timer** field.

Recovery:

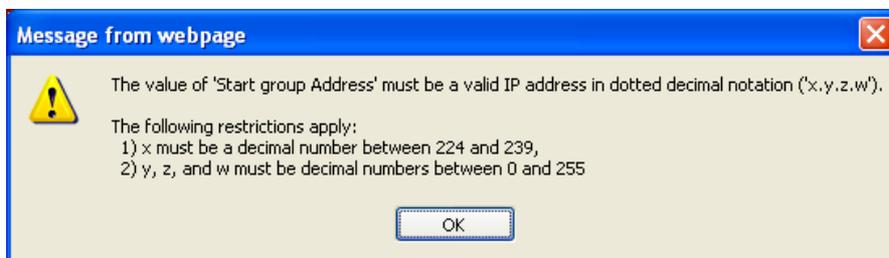
1. Click the **OK** button to clear the webpage message.
2. Make the requested configuration change.
3. See the '[MRP Configuration](#)' section.

Message:

The value of 'add group address' must be a valid IP address in dotted decimal notation ('x.y.z.w').
 The value of 'Start group Address' must be a valid IP address in dotted decimal notation ('x.y.z.w').

The following restrictions apply:

- 1) x must be a decimal number between 224 and 239,
- 2) y, z, and w must be decimal numbers between 0 and 255



Meaning: At **Monitor > MVR > Groups Information** you entered an invalid entry in the “**add group address**” field.

Recovery:

1. At **Monitor > MVR > Groups Information** enter a valid entry in the “add group address” field.
2. See the “[Monitor > MVR > Groups Information](#)” section.

Problem: Lost management of LIB-4400

LIB-4400 receiving excess broadcast packets causes loss of management

Meaning: High amount of broadcast packets received on a port will cause loss of management of the LIB-44xx. This happens, even when STP is blocking the port which is the source of broadcasts, and when QoS Storm Control feature is configured to limit broadcast packets as low as 1pps.

LIB-44xx Configuration: STP is enabled (default settings) and QoS Storm Control is enabled (broadcast, multicast, and unicast all limited to low value such as 1pps).

After connecting the LIB-44xx to the broadcast storm, it's usually still manageable for several seconds. During this time period before loss of management, observe the following:

STP usually changes the storming port state to blocking.

CPU utilization goes to 99-100% (even when the storming port is discarding).

Port Statistics show massive amounts of packets received on the storm port (even when the storming port is discarding).

Eventually management of the LIB-44xx is completely lost, until it's physically disconnected from the storm. At times, the LIB-44xx is unresponsive even after disconnecting it from the storm, and has to be power cycled to regain management. Disabling Storm Control produces the same results.

Note: when the LIB-44xx is connected to the storming switch, the storm appears to remain confined to just one LIB-44xx port; storm traffic does not spread to other LIB-44xx ports.

Tests through the LIB-44xx are successful and indicate line rate transfers, even while the LIB-44xx is connected to the storm, and is un-manageable.

Recovery:

1. Contact Tech Support.

Message: **There isn't any entry provide WEB service. Do you want to proceed anyway?**



Meaning: At **Configuration > Security > Switch > Access Management** you enabled Access Management Configuration and clicked the Save button without first adding a new entry.

Recovery:

1. Click the **Cancel** button and add a new entry, or click the **OK** button and continue operation.
2. See “[DHCP Configuration](#)” on page 91.

Message: **Please make sure the DHCP server connected on trust port?**



Meaning: At **Configuration > Security > Network > DHCP > Relay** when you enabled Relay Mode and entered a Relay Server IP address, the system requires a DHCP server on a trusted port at the specified IP address.

Recovery:

1. Make sure the DHCP server is connected on a trust port.
2. See “[DHCP Configuration](#)” on page 91.

**Message: Board Type not found, probing enabled
Invalid configuration detected (Signature Check Failed)**

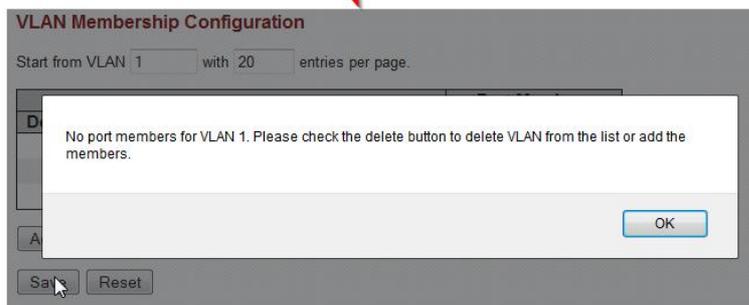
Meaning: The LIB-44xx board type parameter could not be discovered.

Recovery:

1. Make sure the Device ID entry is valid.
2. If possible, reset the LIB-44xx to factory defaults.
3. If possible, reboot the LIB-44xx.

Message: No port members for VLAN 1. Please check the delete button to delete VLAN from the list or add the members.

Delete	VLAN ID	VLAN Name	Port Members										
			1	2	3	4	5	6	7	8			
<input type="checkbox"/>	1	default	<input type="checkbox"/>										
<input type="checkbox"/>	2		<input checked="" type="checkbox"/>										
<input type="checkbox"/>	3		<input checked="" type="checkbox"/>										



Meaning: You tried to delete a non- existent VLAN translation entry from a group.

Recovery:

1. Click the **OK** button to clear the message.
2. Re-enter the command with a different (existing) group.
2. Add port member to the group and re-try the operation.
3. Make sure you are not trying to delete VLAN 1. Deleting VLAN 1 causes issues with forwarding.
4. To make sure no ports belong to VLAN 1, and then add VLAN 1 with all ports set to the forbidden state.
5. See the **Delete VLAN Translation Group Entry** command for CLI information.

Problem: **CPU can be overloaded by broadcast packets, causing loss of management.****LIB-4400 receiving excess broadcast packets causes loss of management.**

Meaning: LIB-44xx CPU overload and loss of management occur even when STP is discarding all packets from the port that the broadcasts are from, and when QoS Storm Control is configured to limit broadcast packets to rates as low as 1pps.

Recovery:

1. Enable Loop Protection.
2. See "[Loop Protection Configuration](#)" on page 117.

Problem: **Ingress Bandwidth profiling on Port doesn't allow for bursty TCP flows.**

Meaning: The Port ingress policers work fine for traffic generation for normal L2 traffic streams. For TCP flow which is bursty in nature, the resultant bandwidth is very poor close to 10% of the set rate. The issue can be that the leaky bucket for port ingress policer may not be set correctly to accommodate for the bursts mainly for TCP control frames.

Recovery:

1. Use EPL service which allows for bursty traffic flow.
2. See "[Ethernet Services Configuration](#)" on page 220.

Problem: **MEP not working over link aggregation.**

Meaning: Protocols above the LAG layer don't seem to consider the LAG group as a logical port (e.g., SOAM over an aggregated link). A MEP does not view the LAG group as a logical port; for example, if you:

1. Create an EVC with all LAG groups' ports in NNI.
2. Create a LAG group.
3. Create MEPs on the EVC (e.g., EVPL) service on UNI and NNI ports. Note that since a MEP has a residence port attached to it even though it's an EVC MEP, this creates issue when the residence port is down and CCMs are to be carried over the other ports. This creates faults on the MEP, possibly because of the MAC address being used in the CCMs.

Recovery: This is a deployment issue.

1. Use MEP on EVC instead; this is orthogonal to any aggregation. For a service running over an aggregation, just add all port to the NNI.

Problem: **DNS not updated when new DHCP address is granted**

Meaning: When a new address is granted a device via a DHCP operation, the "A" and "PTR" records in DNS must be corrected to point at the new address. The Client (3280) should drive that by sending "Option 81" in with the DHCP request response. This doesn't appear to be happening on

the 3280 as a ping to the DNS name will fail. The 'A' record maps a host name to an IP address and the 'PTR' record creates a pointer to the host for reverse lookups.

Recovery: FQDN option 81 refers to the Fully Qualified Domain Name (FQDN) Dynamic Host Configuration Protocol (DHCP) option (81). See Microsoft TechNet article # bb727018 at <http://technet.microsoft.com/en-us/library/bb727018.aspx>.

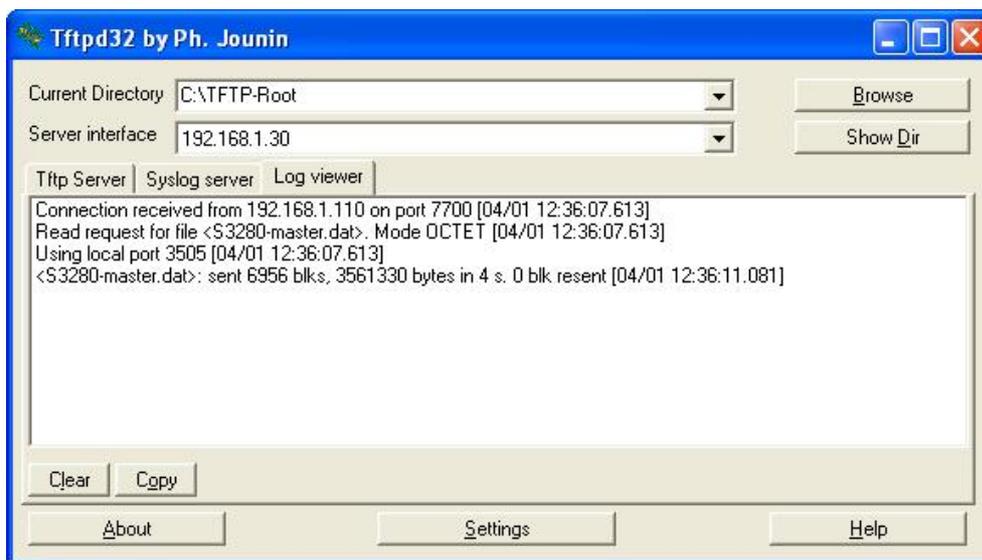
Message:

Connection received from 192.168.1.110 on port 7700 [04/01 12:36:07.613]

Read request for file <LIB-4400-master.dat>. Mode OCTET [04/01 12:36:07.613]

Using local port 3505 [04/01 12:36:07.613]

<LIB-4400-master.dat>: sent 6956 blks, 3561330 bytes in 4 s. 0 blk resent [04/01 12:36:11.081]



Meaning: Standard messages received from TFTPd32 Log Viewer tab.

Recovery: None; information only. Tftpd32 is an open source IPv6 ready application with free DHCP, TFTP, DNS, SNTP and Syslog servers, and a TFTP client. The TFTP client and server are compatible with TFTP option support (e.g., tsize, blocksize, and timeout). Some extended features (e.g., directory facility, security tuning, interface filtering; progress bars and early acknowledgments) enhance the TFTP protocol usability and transfer rate for both client and server. The included DHCP server provides unlimited automatic or static IP address assignment. Tftpd32 is also provided as a Windows service. Tftpd64 is the same application compiled as a 64 bits application. See the TFTPd32 FAQ at http://tftpd32.jounin.net/tftpd32_faq.html#static_DHCP_english.

Message: Peer MEP ID x is already in use

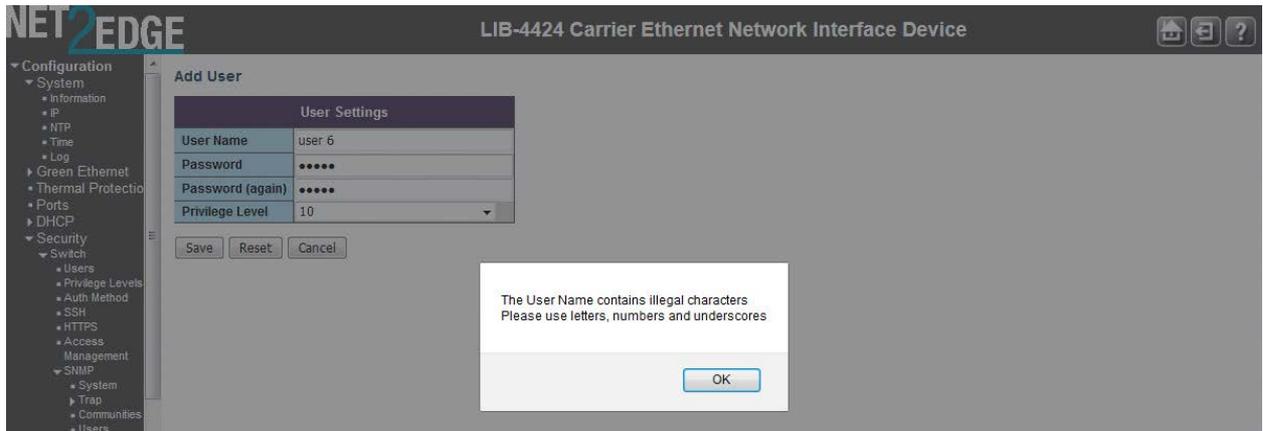
The screenshot shows the NET2EDGE web interface for a LIB-4424 Carrier Ethernet Network Interface Device. A modal dialog box is displayed in the center with the message "Peer MEP ID 2 is already in use" and an "OK" button. The background configuration page shows the "Peer MEP Configuration" section with a table containing one entry with Peer MEP ID 2 and Unicast Peer MAC 00-00-00-00-00-00. Other sections like "Instance Configuration", "Functional Configuration", and "TLV Configuration" are also visible.

Meaning: At the **Configuration > MEP > MEP Configuration** menu path you tried to add a new peer MEP with an existing Peer MEP ID.

Recovery:

1. Click the **OK** button to close the webpage message.
2. Enter a new Peer MEP ID and click the **Save** button.
3. See "[MEP Instance Configuration](#)" on page 162.

Message: The user name contains illegal characters Please use a combination of letters, numbers and underscores.

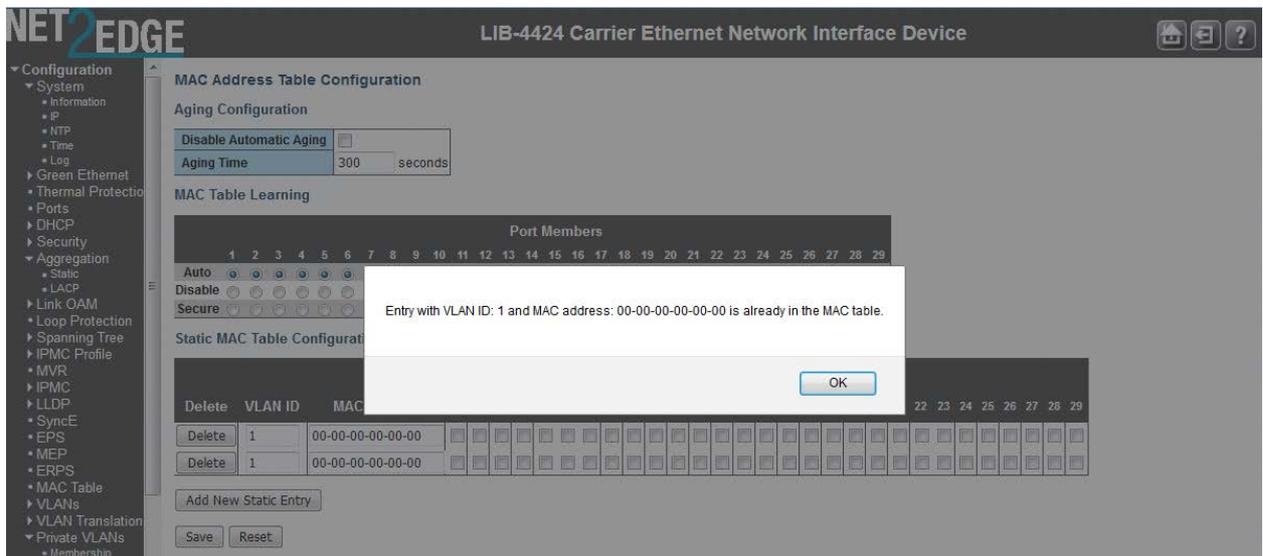


Meaning: At **Configuration > Security > Switch > Users** you entered an unacceptable character in the User Name field.

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Enter a new User Name without any spaces or other illegal characters and click the **Save** button.
4. See the “[Add a New User](#)” section on page 30.

Message: Entry with VLAN ID: x and MAC Address xx-xx-xx-xx-xx-xx already in MAC table.



Meaning: At the **Configuration > MAC Table > Static MAC Table Configuration** menu path, you tried to add a new static entry with a VLAN ID and MAC address that was already used.

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Enter a new static entry with a VLAN ID and MAC Address that is not already used.

3. Click the **Save** button.
4. See “[MAC Address Table Configuration](#)” on page 197.

Problem: Management Port EtherType 9100 does not function

Meaning: When Management Port Ethertype Customer S-port is set to 9100, LIB-44xx access via Port x (MGMT port) is lost. This is a known issue with a fix in process, available at the next LIB-44xx version release.

Recovery:

1. Make sure you are running the latest version of LIB-44xx software; upgrade if a newer version is available.
2. If possible, use the LIB-44xx CONSOLE PORT.
3. If possible, use Ethertype 88A8 or 8100. See the “[Ports](#)” description on page 225.
4. Verify the APS and MEP configurations.
5. Retry the operation.

Problem: When ERPS is configured using a separate APS and SF MEPs, the LIB-4400 will crash.

The separate APS MEP is configured with CCM disabled and APS enabled.

Meaning: This is a known issue with a fix in process; available at the next LIB-44xx version release.

The LIB-44xx boot script runs, and a series of messages displays:

```
Warning: conf_sec_open failed or size mismatch, creating defaultsW erps 00:00:01
29/erps_init#1315: Warning: conf_sec_open failed or size
mismatch, creating defaultsW link_oam 00:00:01 29/eth_link_oam_init#3012: Warning:
conf_sec_open failed or size mismatch, creating
defaultsPassword:
Login in progress...
Invalid username or password!
Username: adminPassword:
Login in progress...
Welcome to Net2Edge. Command Line Interface (v1.0).
Type 'help' or '?' to get help.
```

Recovery:

1. Check the IP configuration. At the CLI prompt type **ip conf'** and press **Enter**. For example:

```
>ip conf
IP Configuration:
=====
```

```
DHCP Client   : Enabled
IP Address    : 192.0.2.1
IP Mask       : 255.255.255.0
IP Router     : 0.0.0.0
VLAN ID       : 1
```

3. Make sure you are running the latest version of LIB-44xx software; upgrade if a newer version is available.
4. Verify the APS and MEP configurations. See the related sections of this manual.
5. Retry the operation.

Message: **fis load fails after firmware upgrade.** The 'fis load -d managed' fails with one of the following errors after firmware upgrade:

decompression error: invalid block type

decompression error: invalid stored block lengths

Meaning: Flash Corruption and SPI Bus Access bug.

Recovery:

1. Boot managed.bk and run the firmware upgrade again.
2. When running the firmware upgrade again does not work, reprogramming the flash image has worked. Contact Support for direction.

Problem: **Auto Negotiation between a 2.5G port and a 1G port does not work in LIB-4400 v1.0.**

Meaning: A check was added '&& (conf->speed == VTSS_SPEED_1G)', but 2.5G was missed.

The LIB-44xx software reads the SFP module and configures the link accordingly. On some systems it is not possible to read the SFP module via the I2C interface and this is the reason for the faulty behaviour.

This is a known issue with a fix in process; available at the next LIB-44xx version release.

Recovery:

1. Make sure you are running the latest version of LIB-44xx software; upgrade if a newer version is available.
2. Verify the port and auto-negotiation configurations. See the related sections of this manual.
3. Retry the operation.

Problem: **Security level 11 appears to be equivalent to level 1**

Meaning: When running at security level 11, only commands that work at levels 1 - 4 will display via the help, and do not seem to be otherwise available. This is a known issue with a fix in process, available at the next LIB-44xx version release.

Recovery:

1. Do not use Security Level 11. Use Levels 10 and 12-14 instead.

Problem: **Security level 1 help does not display correct results.**

Meaning: Entering "S ?" in the CLI running under level 1 displays commands that do not work in level 1, and do not begin with the letter "s". The "System" command is available in limited form, but to get that information, "sy ?" must be entered. This is a known issue with a fix in process, available at the next LIB-44xx version release.

Recovery: Use the "sy ?" command instead of the "S ?" command. For example:

1. Enter an "S ?" command which displays:

Command Groups:

Switch : Switch security
 Network : Network security
 AAA : Authentication, Authorization and Accounting

Type '<group>' to enter command group
 Type '<group> ?' to get group help

2. Enter an "**sy ?**" command which displays:

Available Commands:

```
System Configuration [all | (port <port_list>)]
System Log Configuration
System Version
System Log Lookup [<log_id>] [all|info|warning|error]
```

Problem: Errors displaying Sys Log Lookup command data.

Meaning: A series of cli_parser command errors display with the "sys log lookup" command. For example:

sys log lookup

Number of entries:

Info : 5

Warning: 0

Error : 176

All : 181

ID Level Time Message

```
-----
-
1 Info - Switch just made a cool boot.
2 Info 1970-01-01T00:00:02+00:00 Using primary power source.
3 Info 1970-01-01T00:00:09+00:00 Link up on port 1
4 Info 1970-01-04T02:51:12+00:00 Link down on port 1
5 Info 1970-01-04T02:51:46+00:00 Link up on port 1
6 Error 1970-01-07T00:08:56+00:00 E cli 97/cli_parse_command#2864: Er ...
7 Error 1970-01-07T00:08:56+00:00 E cli 97/cli_parse_command#2879: Er ...
8 Error 1970-01-07T00:08:56+00:00 E cli 97/cli_parse_command#2879: Er ...
9 Error 1970-01-07T00:08:56+00:00 E cli 97/cli_parse_command#2879: Er ...
```

Recovery:

1. Make sure you are running the latest version of LIB-44xx software; upgrade if a newer version is available.
2. If possible, use the existing information as displayed.

Problem: Security levels 5 - 9 show strange results when **n ?** is entered.

Meaning: Entering **N ?** in other security levels will show the NULL command, but when done in levels 5 - 9, writes several lines of NULL. This is a known issue with a fix in process, available at the next LIB-44xx version release. For example:

```
test1:/>NULL ?
```

Available Commands:

```
NULL
```

```

NULL
test1:/>

```

Recovery:

1. Make sure you are running the latest version of LIB-44xx software; upgrade if a newer version is available.
2. If possible, use the existing information displayed.

Problem: the input **n ? ?** displays internal errors.

Meaning: Entering the command "**n > ?**" causes 'parse_command' internal error 2879 to display repeatedly. For example:

```

test1:/>n ? ?
E cli 139/cli_parse_command#2864: Error: Internal error
E cli 139/cli_parse_command#2879: Error: Internal error

```

Recovery:

1. Make sure you are running the latest version of LIB-44xx software; upgrade if a newer version is available.
2. If possible, use the existing information displayed, or try using another CLI command.

Problem: Can no longer HTTPS browse to LIB-44xx after new certificate is generated

Meaning: Initially enabling HTTPS and browsing to the LIB-44xx with HTTPS works. However, if a new certificate is generated, any web browser that previously navigated the LIB-44xx via https (using

the old certificate) can no longer HTTPS browse the LIB-44xx. This may be because when a new certificate is generated, it re-uses the original certificate's serial number. This is a known issue with a fix in process, available at the next LIB-44xx version release.

Recovery:

1. Make sure you are running the latest version of LIB-44xx software; upgrade if a newer version is available.
2. If possible, use the existing certificate.

Problem: The **Clear** button does not clear data.

Meaning: Using the **Clear** button does not clear data at **Monitor > Link OAM > Event Status**. This is a known issue with a fix in process, available at the next LIB-44xx version release.

Recovery:

1. Make sure you are running the latest version of LIB-44xx software; upgrade if a newer version is available.
2. Try using the **Refresh** and/or **Auto-refresh** buttons.
3. Try using the **Clear** button at **Monitor > Link OAM > Statistics**.
4. Adjust the configuration at **Configuration > Link OAM > Port Settings** or at **Configuration > Link OAM > Event Settings**.

Problem: The **Clear** button at **Configuration > Link OAM > Event Settings** causes issues.

Meaning: Clicking the **Clear** button from this menu path will reset to all zeros in Error Window and Error Threshold, but refreshing will display default values again and clear the data at **Monitor > Link OAM > Statistics**. This is a known issue with a fix in process, available at the next LIB-44xx version release.

Recovery:

1. Make sure you are running the latest version of LIB-44xx software; upgrade if a newer version is available.
2. If possible, do not use the **Clear** button.

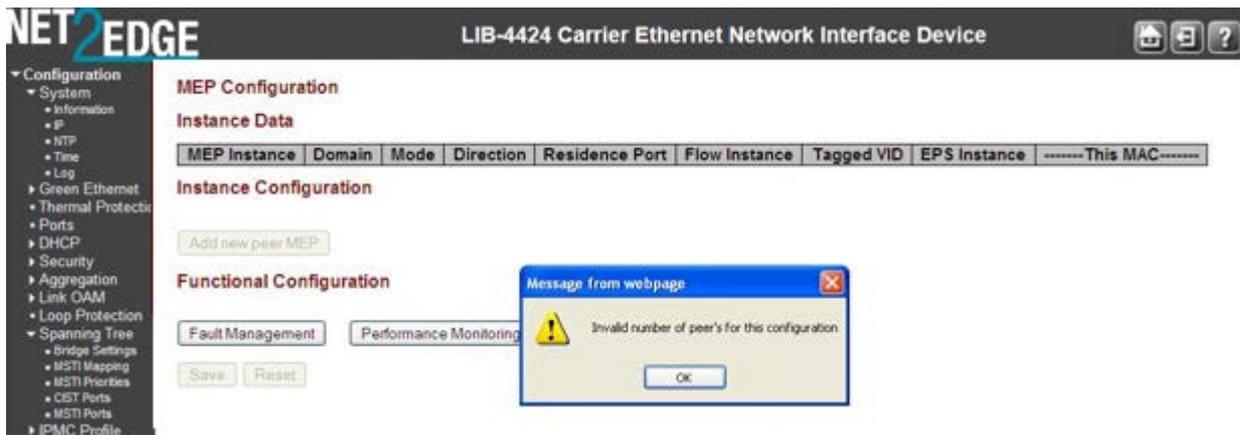
Problem: **Port admin status shows as 'disabled'**.

Meaning: The Port Admin status always displays as 'disabled', even when it is actually enabled. This is a known issue with a fix in process, available at the next LIB-44xx version release.

Recovery:

1. Make sure you are running the latest version of LIB-44xx software; upgrade if a newer version is available.
2. If possible, ignore the status displayed.

Message: **Invalid number of peer's for this configuration**

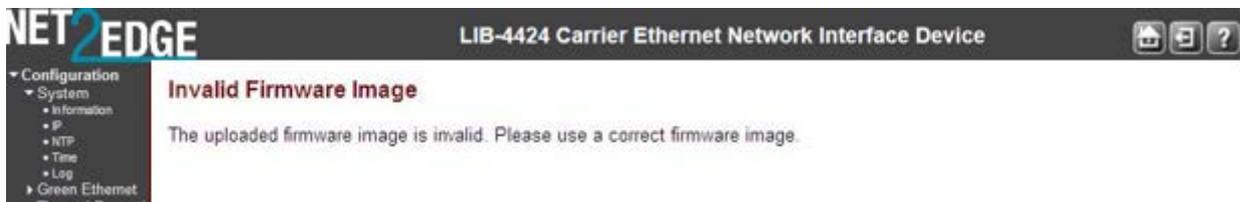


Meaning: At **Configuration > MEP > MEP Configuration > Instance Configuration** > you clicked the **Save** button before adding any Peer MEPS.

Recovery:

1. Click the **OK** button to clear the webpage message.
2. At **Configuration > MEP > MEP Configuration > Instance Configuration** > add one or more valid Peer MEPS and then click the **Save** button.
3. See “[MEP Instance Configuration](#)” on page 162.

Message: **Invalid Firmware Image - The uploaded firmware image is invalid. Please use a correct firmware image.**



Meaning: At **Maintenance > Software > Upload** you clicked the **Upload** button without first browsing to and selecting a valid LIB-44xx firmware filename.

Recovery:

1. Click the browser **Back** button to return to the Firmware Update page.
2. Browse to and select a valid file to use for this upgrade (e.g., *LIB-4400-v1.0.1.dat*).
3. At the ‘Choose File to Upload’ dialog box, click the **Open** button.
4. Click the **Upload** button.
5. At the web page message, select **OK** or **Cancel**.
6. The upgrade will proceed, or the message “*Firmware Upload Error - Flash is already updated with this image*” will display.
7. Refer to the “[Software Upload Procedure](#)” on page 420.

Message: **Firmware Upload Error - Flash is already updated with this image**



Meaning: You tried to update the LIB-44xx with the current (existing) firmware version.

Recovery:

1. Click the browser **Back** button to return to the Firmware Update page.
2. Browse to and select a valid file to use for this upgrade (e.g., *LIB-4400-v1.0.1.dat*).
3. At the 'Choose File to Upload' dialog box, click the **Open** button.
4. Click the **Upload** button.
5. At the web page message, select **OK** or **Cancel**.
6. The upgrade will proceed, or the message "*Firmware Upload Error - Flash is already updated with this image*" will display.
7. Refer to the "[Software Upload Procedure](#)" on page 420.

Message: **Warning! Device will automatically reboot. Proceed with update now?**



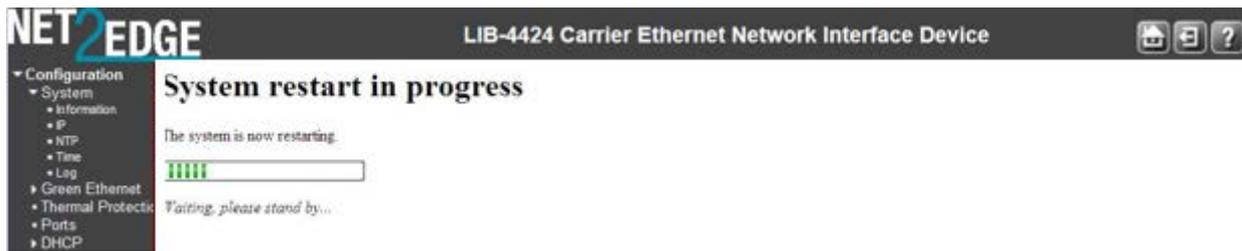
Meaning: Information message indicating that you initiated a firmware update from the **Maintenance > Software > Upload** menu path. This procedure transfers the uploaded firmware image to the LIB-44xx flash component. The LIB-44xx will restart after the update.

Note: do not reset or power off the LIB-44xx during this process.

Recovery:

1. Verify that you want to perform this firmware update and resultant reboot and click **OK**. Otherwise click **Cancel**.
2. Refer to the "[Software Upload Procedure](#)" on page 420.

Problem: The message "*System restart in progress - The system is now restarting. - Polling...*" displays continuously without completing.



Meaning: At **Maintenance > Restart Device > Are you sure you want to perform a Restart?** prompt, you selected **Yes** with the **Force Cool Restart** checkbox checked.

Recovery:

1. Press the keyboard **Esc** key.
2. Press the browser **Back** button.
3. Log out and then log back in to the LIB-44xx.
4. Log out of the LIB-44xx, close the browser session, and then log back in to the LIB-44xx.
5. Restart the LIB-44xx from the **Maintenance > Restart Device** menu path.

Message:

error in putting all port+vlan in forwarding mode
failing in putting blocked port in forwarding state
Error in enabling forwarding for group = erpg->group_id

Meaning: Protection switching can be triggered by fault conditions and external manual commands. The fault conditions include Signal Failure (SF) and Signal Degrade (SD), where:

sf_state : signal failure state on a given ring port

sd_state : signal degrade state of a given ring port (for future use)

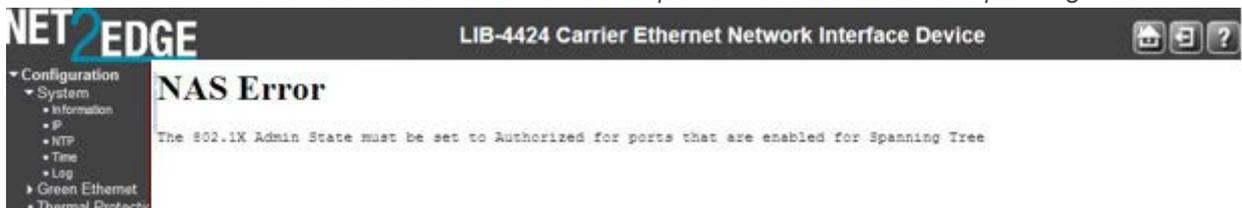
An SD fault occurs when at least one of the links it forwarded through has rate of errored bits exceeding a predetermined BER threshold.

The default configuration is that all VLANs are disabled and all ports are discarding for all ERPS instances.

Recovery: If the application (ERPS module) enables a VLAN for an ERPS instance, it must set up the forwarding state for all ports for the instance.

Message: **NAS Error**

The 802.1X Admin State must be set to Authorized for ports that are enabled for Spanning Tree



Meaning: At **Configuration > Security > Network > NAS >** menu path you selected an unsupported NAS configuration for the existing Spanning Tree configuration.

Recovery:

1. Click the browser Back button to return to the **Configuration > Security > Network > NAS >** page.
- 2a. In the **Port Configuration** section in the **Admin State** column, select **Force Authorized**.
or:
- 2b. At **Configuration > Spanning Tree > CIST Ports** in the CIST Normal Port Configuration table in the STP Enabled column, uncheck the checkbox for the ports you want to set to other than **Force Authorized**.
3. Click the **Save** button.
4. At **Configuration > Security > Network > NAS >** menu path, in the **Port Configuration** section in the **Admin State** column, check the checkbox for the ports you want to set to other than Force Authorized (i.e., select Force Unauthorized, Port-based 802.1X, Single 802.1X, Multi 802.1X, or MAC-based Auth.).

**Message: >E api 00:28:25 57/vtss_vcap_add#640: Error: VCAP IS1: Could not find ID: 2
E web 00:28:25 57/handler_config_evc_edit#311: Error: evc_mgmt_add(0): failed**

Meaning: At the LIB-44xx **Configuration > Ethernet Services** menu path, you tried to set up the ECE before setting up the related EVC.

At the CLI, you entered one of the EVC ECE commands before entering the related EVC commands.

Recovery: Set up the EVC before trying to set up the ECE.

Message: The new setting may lost some dynamic entries of port 2. Do you want to proceed anyway?

Meaning: At **Configuration > Security > Network > IP Source Guard > Static Table** you set a port configuration and clicked the **Save** button.

Recovery:

1. Click the **OK** button to continue or click **Cancel** to stop.
2. Make sure that this is the port configuration that you want. See **Configuration > Security > Network > IP Source Guard > Configuration** for more information.

Message: Switch does not respond.

There was a problem getting page data: Unknown

Meaning: One of the webpage messages above displays and the **Translate dynamic to static** button and the **Save** and **Reset** buttons are greyed out at the **Configuration > Security > Network > IP Source Guard > Configuration** page.

Recovery:

1. Click the **OK** button to clear the webpage message.
2. Re-enter the parameter within the valid range. See "[IP Source Guard Configuration](#)" on page 95.
3. Click the **Save** button when done.
4. Try the CLI command "**config default keep**" to restore the default setting but keep the assigned IP address, and then re-try the IP source guard operation.

Problem: When the FDB table reaches 8192 entries, a new dynamic MAC will be learned and override the old one.

Description: When the MAC table is full (8192 entries), the switch learns new entries and purges older entries even though these entries are not aged. This is the expected behaviour. The MAC table is implemented so that it will always store the most active entries. The least recently used entries are removed to provide storage for new entries. The benefit of this approach is that in a normal network, flooding is kept to a minimum when the MAC table becomes full. This occurs even when aging is disabled. This behaviour is not configurable.

Resolution:

1. Apply Port- or MAC-based authentication (IEEE802.1X).
2. Set up a per-port MAC table limit, preventing an intruder from taking up too much of the MAC table. For example, if a port is limited to 64 entries in the MAC table, frames causing this limit to be exceeded are discarded, and the MAC table is not affected.
3. Set up additional Storm Policers to prevent flooding on an unknown MAC address. See the **Configuration > QoS** section.

Problem: The "psec_limit" cannot work when configuration "limit" is "1024".

The LIB-44xx fails to send the trap only when the limit control reaches '1024' at the 5 packets/sec. rate.

Also, some the highwarn and normal DMI traps occur when the limit parameter reaches '1023' or '1024'.

Description: This is a Port Security Limit Control issue where the software is unable to process the new Mac address frames in time. It is receiving frames faster than it can add to the Mac table and in turn looks like it is losing frames, thus not reaching the limit. In a normal scenario if new Mac addresses turns up for example, once a second, it will work. But in case of an attack of new Mac address, then the port won't shutdown, but it won't switch these frames either. The LIB-44xx works fine when tested at 5 frames/second.

Example:

Port Security Limit Control Configuration

System Configuration

Mode	Enabled
Aging Enabled	<input checked="" type="checkbox"/>
Aging Period	3600 seconds

Port Configuration

Port	Mode	Limit	Action	State	Re-open
*	<>	1024	<>		
1	Enabled	1024	Trap & Shutdown	Ready	Reopen
2	Enabled	1024	Trap & Shutdown	Ready	Reopen
3	Enabled	1024	Trap & Shutdown	Ready	Reopen
4	Enabled	1024	Trap & Shutdown	Ready	Reopen
5	Enabled	1024	Trap & Shutdown	Ready	Reopen
6	Enabled	1024	Trap & Shutdown	Ready	Reopen
7	Enabled	1024	Trap & Shutdown	Ready	Reopen
8	Enabled	1024	Trap & Shutdown	Ready	Reopen
9	Enabled	100	Trap	Ready	Reopen
10	Disabled	4	None	Disabled	Reopen

Save Reset

SNMP System Configuration

Mode	Enabled
Version	SNMP v2c
Read Community	public
Write Community	private
Engine ID	800007e5017f000001

SNMP Trap Configuration

Trap Mode	Enabled
Trap Version	SNMP v2c
Trap Community	public
Trap Destination Address	192.168.0.37
Trap Destination IPv6 Address	::
Trap Authentication Failure	Enabled
Trap Link-up and Link-down	Enabled
Trap Inform Mode	Enabled
Trap Inform Timeout (seconds)	1
Trap Inform Retry Times	5

Save Reset

Port Configuration

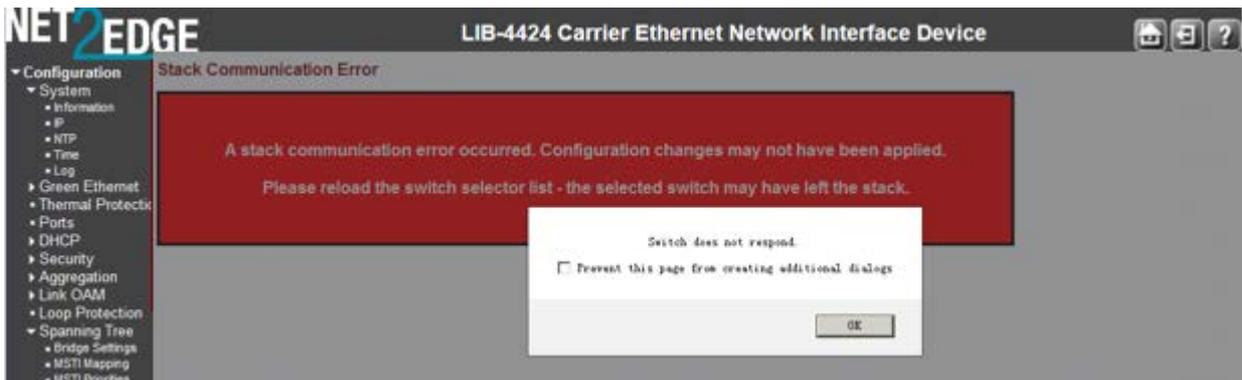
Port	Link	Speed		Flow Control			Maximum Frame Size	Excessive Collision Mode	Power Control	Description
		Current	Configured	Current Rx	Current Tx	Configured				
*		<>	<>			<input type="checkbox"/>	9600	<>	<>	
1	● 100fdx	Auto	Auto	×	×	<input type="checkbox"/>	9600	Discard	Disabled	
2	● 100fdx	Auto	Auto	×	×	<input type="checkbox"/>	9600	Discard	Disabled	
3	● 100fdx	Auto	Auto	×	×	<input type="checkbox"/>	9600	Discard	Disabled	
4	● 100fdx	Auto	Auto	×	×	<input type="checkbox"/>	9600	Discard	Disabled	
5	● 1Gfdx	Auto	Auto	×	×	<input type="checkbox"/>	9600			
6	● 1Gfdx	Auto	Auto	×	×	<input type="checkbox"/>	9600			
7	● 1Gfdx	Auto	Auto	×	×	<input type="checkbox"/>	9600			
8	● 1Gfdx	Auto	Auto	×	×	<input type="checkbox"/>	9600			
9	● 100fdx	Auto	Auto	×	×	<input type="checkbox"/>	1518	Discard	Disabled	

Save Reset

Resolution:

Decrease the limit to less than 200 (e.g., **security network limit 194**).
 Check for upgrades with a fix to this known problem.

Message: Stack Communication Error - A stack communication error occurred. Configuration changes may not have been applied. Please reload the switch selector list - the selected switch may have left the stack. Switch does not respond. **Prevent this page from creating additional dialogs** - OK



Meaning: An internal error occurred during Flow Control configuration at the **Configuration > Port > Configuration** menu path. The 'current Rx/Tx' seems to be Link partner's ability for pause.

Port	Link	Speed		Flow Control			Maximum Frame Size
		Current	Configured	Current Rx	Current Tx	Configured	
*		<>	<>			<input type="checkbox"/>	9600
1	● Down	Auto	Auto	×	×	<input type="checkbox"/>	9600
2	● Down	Auto	Auto	×	×	<input type="checkbox"/>	9600
3	● 10fdx	10Mbps FDX	10Mbps FDX	×	×	<input checked="" type="checkbox"/>	9600
4	● 100fdx	100Mbps FDX	100Mbps FDX	×	×	<input checked="" type="checkbox"/>	9600
5	● Down	Auto	Auto	×	×	<input type="checkbox"/>	9600
6	● Down	Auto	Auto	×	×	<input type="checkbox"/>	9600
7	● Down	Auto	Auto	×	×	<input type="checkbox"/>	9600
8	● Down	Auto	Auto	×	×	<input type="checkbox"/>	9600
9	● 100fdx	Auto	Auto	×	×	<input type="checkbox"/>	1518

Recovery:

1. Check or uncheck the checkbox.
2. Click the **OK** button.

Message: The Connection has timed out

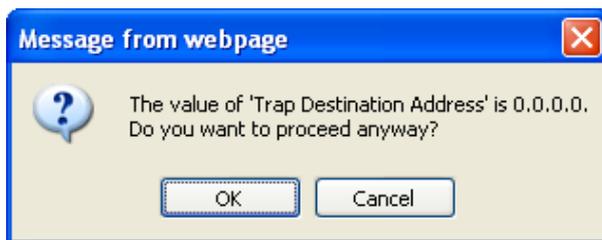


Meaning: The LIB-44xx switch cannot communicate with the PC after enabling Flow Control from the **Configuration > Port > Configuration** menu path.

Recovery:

Change the MGMT port to 1522 as the default frame size. Increasing the CONSOLE PORT Max Frame size increased to tagged frame length fixes the problem.

Message: **The value of 'Trap Destination Address' is 0.0.0.0. Do you want to proceed anyway?**



Meaning: At the **Configuration > Security > Switch > SNMP > System** menu path you clicked the Save button without making an entry in the “Trap Destination Address” field.

Recovery:

1. If you want to enter a Trap Destination Address, click the **Cancel** button and make the entry. If you want to proceed anyway, click the **OK** button.

Problem: **'System Restore Default' hangs intermittently and indefinitely.**

Meaning: This appears to only happen when an EVC/ECE is configured. In LIB-44xx v1.2.1 and above, EVC and ECE configurations are removed when you execute a System Restore via the web (**Maintenance > Configuration > Restore Binary**) or via the CLI (**config restore binary** command).

Recovery:

1. If the EVC/ECE is removed before defaulting then the command will not hang.
2. Upgrade the LIB-44xx to the latest software version.

Problem: At **Configuration > Ethernet Services > Bandwidth Profiles** in the “Bandwidth Profiles Configuration” table, the numbers in the CIR (kbps) CBS (bytes) EIR (kbps) EBS (bytes) columns are different than expected.

Meaning: EVC BWP rates are now calculated from layer 2. Software versions 1.2.2 and before are based on Line rate (level 1) in BWP; versions 1.2.3 and above are based on Data rate (level 2) in BWP (per MEF CE2.0).

Recovery:

1. Reconfigure the “Bandwidth Profiles Configuration” table parameters.
2. See [“Configuration > Ethernet Services > Bandwidth Profiles”](#) on page 224.

Message:

*Peer MAC is needed for HW based CCM.
peer MAC address should be unicast
Peer MAC should be unicast.*

Meaning: The "Unicast Peer MAC" option only supports the unicast address to be configured.

Recovery:

1. Change the MEP configuration. See the “Peer MEP Configuration” section (**Configuration > MEP > MEP Configuration** menu path).
2. Retry the operation.

Message: *Setting Tx mirroring for mirror port (port %lu) has no effect. Tx mirroring ignored.*

Meaning: At **Configuration > Mirroring > Mirror Port Configuration** table, if you set an invalid value to the Mirroring Mode, this SNMP error displays.

Recovery:

1. Select a valid Mirror Mode entry (e.g., enable, disable, rx, or tx).
2. See [“Mirroring Configuration”](#) on page 316.

Message:

*syslog Clear Level = %ld
syslog message get info fail!
syslog server address to set is: %s
Testing syslog number %d (prefix: Debug, Info, Warning, etc.)*

Meaning: You entered an invalid Syslog entry.

Recovery:

1. Select a valid Syslog entry.
2. See [“Log \(System Log\) Configuration”](#) on page 32.

Message:

txt2ipv4 failed for tnIpAddr = %s (IP Mgmt Table Entry)
txt2ipv4 failed for tnSubnetMask = %s(IP Mgmt Table Entry)
Configuration failed (DefaultGateway table entry)

Meaning: An error occurred when configuring the DNS Server table.

Recovery:

1. Select a valid DNS entry at **Configuration > System > IP**.
2. See [“IPv4 Configuration”](#) on page 12.

Message: >E api/cil 00:05:10 28/I26_action_check#5539: Error: ACL policer and EVC policer cannot both be enabled

Meaning: Two means of policers cannot be enabled together.

Recovery:

1. Disable one or the other.

Message:

Error: Part of the VLAN configuration didn't apply as some of ports are in forbidden list for a given VLAN id.

Error: Part of the VLAN configuration didn't apply as some of ports are in member list for a given VLAN id.

notice: forbidden port cannot overlap with egress member port in MIB setting.

Meaning: The system does not allow changing VLAN member ports once there is forbidden port (or ports) in this VLAN.

Recovery:

1. Change either the VLAN member ports configuration or change the VLAN forbidden ports configuration.
2. See the **“Configuration > VLANs > VLAN Membership”** section of this manual.

Message:

Warning: FPGA 1.1 required

Warning: FPGA 2.x required

Meaning: The LIB-44xx does not automatically reboot after FPGA upgrade, but should be rebooted to ensure proper operation.

Recovery: We recommend upgrading the FPGA first, then the software, because the unit will automatically reboot after the software upgrade. If upgraded in the reverse order (software then FPGA), you must manually reboot after the FPGA upgrade.

Message: Can't register ACL rule for SAT

Meaning: An ACL function failed (e.g., EtherSAT ACL delete).

Recovery:

1. Retry the operation.
2. See the related section of this manual.

Message:

Error. Invalid value : Frame Size Mix[" + j + "] should be in range [64-9600]

Error. Invalid value : Frame Size Mix[" + j + "] should be in range [" + frame_size_min + "-" + frame_size_max "]

Meaning: A Frame Size error or mismatch occurred.

Recovery:

1. Verify the configured frame size (e.g., 64-9600 or 64-10056).
2. Retry the operation.
3. See the related section of this manual.

Problem: A Reboot will stop Fault Management and Performance Management measurements.

Meaning: After a system reboot, the MEP PM and FM become disabled.

Recovery: 1. Perform a system Reboot.

2. At **Configuration > MEP > MEP Configuration**, re-enable the FM (LB, LT, TST, AIS, LCK) and PM (Loss Measurement and Delay Measurement).

Problem: LIB-44xx reboots at random times.

Meaning: Reboots occur spuriously.

Recovery: 1. Contact support with details. See the "Service" [section](#) on page 512.

Problem: VLAN translation does not appear to work normally when ingress and egress ports are in the same group.

Meaning: If you add VLAN 50 and VLAN 60 to the VLAN table, all the ports are the members of these two VLANs, add LIB-44xx port 2 and port 4 to group 2, set group 2 translate VLAN 50 to VLAN 60, set port 2 and port 4 as C-port, all the others are at default config, continue to send tag=50 frames from IXIA port 1 and port 3, then check both IXIA ports received frames format

Expect result: IXIA port 1 and port 3 should receive tag that has been translated from 50 to 60 frames.

Actual result: All the received frames tags are still 50.

Recovery: None; this is the defined behaviour. When a port belongs to a group, it will do VLAN translation both in ingress IS1 and egress ES0 stage. In the ingress stage, it will do 50 -> 60

translation, while in egress stage, it will do the reverse (60 -> 50) translation. If ingress port and egress port both belong to the same VLAN translation group, it will make the translations counteract each other.

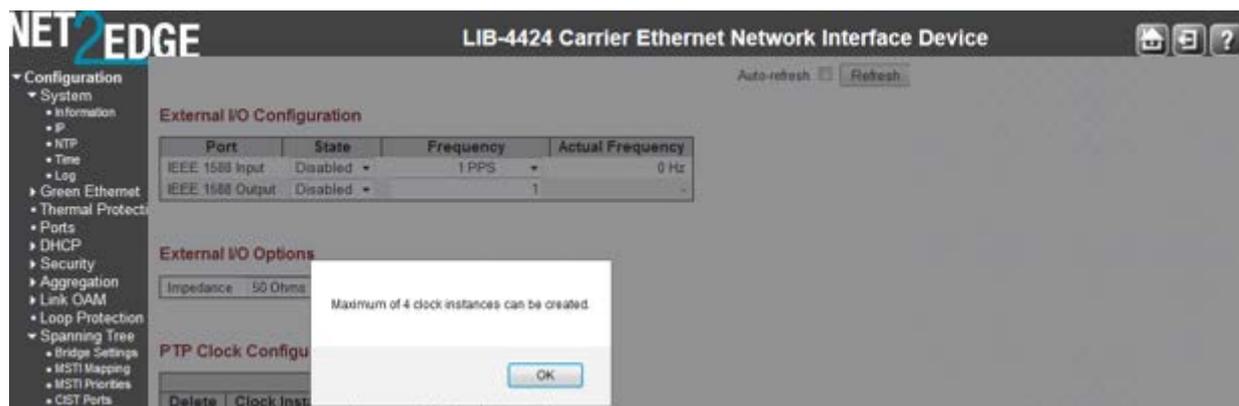
Message: Error. Invalid value : CBS Line Rate should be in range [1-1000] after setting CBS Line Rate to 10000Mbps in LIB-44xx.

Meaning: The RFC2544 CBS Line Rate cannot set to 10000G in LIB-44xx Service Activation Profiles Configuration.

Recovery:

1. Verify the parameter(s) entered. See the related section of this manual.
2. Retry the operation.
3. Contact support with details. See the “[Service](#)” section on page 512.

Message: Maximum of 4 clock instances can be created.



Meaning: You created PTP clock instances 0-2, and then add clock instance 3 as inactive device type, but the PTP clock configuration still shows clock instances 0-2. The “Maximum of four clock instances can be created” window pops up when adding the new PTP clock. In that case no new clock can be added until delete one.

Recovery:

1. Click the **OK** button to clear the message.
2. At **Configuration > SyncE** check the PTP Clock Configuration.
3. At **Configuration > PTP** check the PTP Clock Configuration.

System Log Messages

The LIB-44xx displays four levels of syslog messages as explained below: Note that the **All** level displays all three levels of information that the LIB-44xx can logged (*Info*, *Warning*, and *Error*).

Info Level Messages

These are the Information level messages of the system log. These are normal operational messages used for reporting, measuring throughput, etc. This level of message requires no action.

Table 6: Syslog Info Messages

Info Level Message	Description
<i>Switch just made a cold boot.</i>	The LIB-44xx was restarted. See " Maintenance > Restart Device ".
<i>Link up on port x</i>	The most recent link status on the port x is 'link up'. Port Link is up - no action needed.
<i>Link down on port x</i>	The most recent link status on the port x is 'link down'. See Configuration > Ports > Port Configuration .
<i>Using primary power source.</i>	Normal power on operation.
<i>Frame of 243 bytes received on port 1MAC</i>	Normal frame size information.
<i>Port 1 shut down</i>	Normal port shutdown information.
<i>Power supply 1 present</i> <i>Power supply 2 present</i>	Typical operation.

Warning Level Messages

These messages are the Warning level of the system log. These Warning messages are not an error, but an indication that an error will occur if action is not taken (e.g. *file system 85% full*). Each item must be resolved within a given time.

Table 7: Syslog Warning Messages

Warning Level Message	Description
<i>E api/cil 17:42:26</i> <i>29/126_action_check#6036:</i>	ACL policer and EVC policer cannot both be enabled. Disable one or the other.

Error Level Messages

Error level messages of the system log. These non-urgent failures should be relayed to a developer or an administrator.

Each item must be resolved within a given time.

Table 8: Syslog Error Messages

Error Level Message	Description
<i>E api/cil 17:42:45 29/126_acl_policer_free#6069:</i>	Error: policer 0 already free. The EVC policer or
<i>VLAN Port Configuration Ingress Filter Conflict - MSTP</i>	Verify the Ingress Policers, Port Policing, or Queue Policing configuration. See the Configuration > Security > Network > ACL or the Configuration > QoS menu path.

<i>VLAN Port Configuration Ingress Filter Conflict - ERPS</i>	Verify the VLAN Port and the ERPS configuration. See the related section of this manual.
<i>E web 03:15:49 58/handler_config_vlan#514: Error: vlan_mgmt_vlan_del(1): failed</i>	Verify the Management VLAN configuration. See the related section of this manual.
<i>E link_oam 00:27:32 58/eth_link_oam_mgmt_port_mib_retrival_oper_set#849: Error: Unable to retrieve the mode of the port(1/41)</i>	Verify the Link OAM Mib Retrieve settings. See the “Link OAM Mib Retrieve” section of this manual. Contact Support if necessary; see “ Appendix D: Service, Warranty & Compliance Information ” on page 512.
<i>E ether_sat 00:48:27 62/saDbTestFindEgressPort#2364: Error: SA: Can't find ECE for VID 0, port 1</i>	See “ Service Activation ” on page 310”. See the RFC 2544 User Guide for more information.
<i>E ether_sat 00:48:27 62/handler_config_sa_tests_edit#1440: Error: Error occurred.</i>	See “ Service Activation ” on page 310”. See the EtherSAT User Guide for more information.
<i>E web 00:54:27 62/handler_config_vlan#704: Error: vlan_mgmt_vlan_add(1) sid 1: failed</i>	Verify the Management VLAN configuration. See the related section of this manual.

Note: the **All** level displays all three levels of information that the LIB-44xx can log (*Info, Warning, and Error*).

Third Party Program Messages

The LIB-44xx displays error and information messages from various third party applications, such as Internet Explorer, HyperTerminal, PuTTY, etc. This section lists the messages, provides an example, and discusses the message meaning of and possible recovery steps.

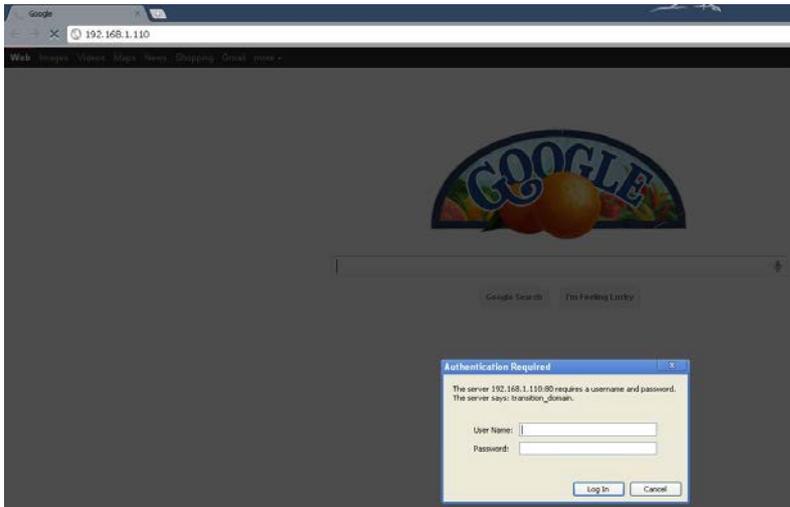
Message: **PuTTY Security Alert - The server’s host key is not cached in the registry.**



Meaning: Normal PuTTY login security message.

Recovery: Click the **Yes** button to trust this host, add the key to the PuTTY cache, and clear the message.

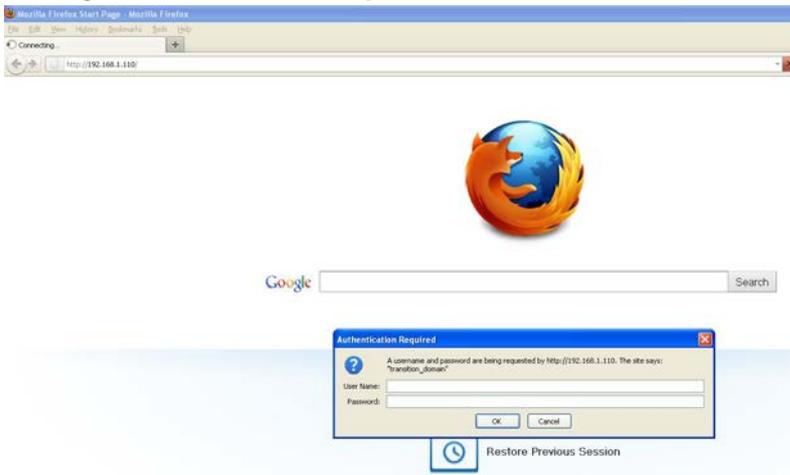
Message: Authentication Required



Meaning: Normal Google Chrome login screen.

Recovery: Enter your LIB-44xx User Name and Password, and click the **Log In** button.

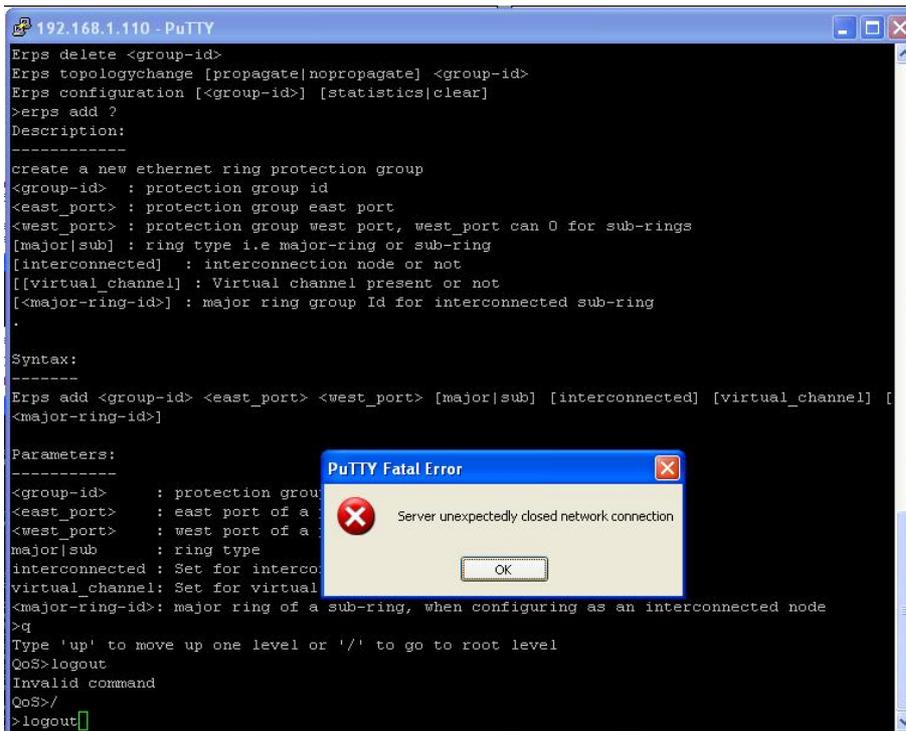
Message: Authentication Required



Meaning: Normal Firefox login screen.

Recovery: Enter your LIB-44xx User Name and Password, and click the **Log In** button.

Message: PuTTY Fatal Error - Server unexpectedly closed network connection



```

192.168.1.110 - PuTTY
Erps delete <group-id>
Erps topologychange [propagate|nopropagate] <group-id>
Erps configuration [<group-id>] [statistics|clear]
>erps add ?
Description:
-----
create a new ethernet ring protection group
<group-id> : protection group id
<east_port> : protection group east port
<west_port> : protection group west port, west_port can 0 for sub-rings
[major|sub] : ring type i.e major-ring or sub-ring
[interconnected] : interconnection node or not
[[virtual_channel] : Virtual channel present or not
[<major-ring-id>] : major ring group Id for interconnected sub-ring
.

Syntax:
-----
Erps add <group-id> <east_port> <west_port> [major|sub] [interconnected] [virtual_channel] [
<major-ring-id>]

Parameters:
-----
<group-id> : protection group id
<east_port> : east port of a protection group
<west_port> : west port of a protection group
major|sub : ring type
interconnected : Set for interconnected node
virtual_channel : Set for virtual channel
<major-ring-id>: major ring of a sub-ring, when configuring as an interconnected node
>q
Type 'up' to move up one level or '/' to go to root level
QoS>logout
Invalid command
QoS>/
>logout

```

Meaning:

Recovery:

1. Click the **OK** button to close the message dialog box.
2. Close the PuTTY session window.
3. Start a new PuTTY session.

Message: Internet Explorer cannot display the webpage



Meaning: The LIB-44xx web interface connection via Microsoft Internet Explorer is down.

Recovery:

1. Try reconnecting via the LIB-44xx web interface in IE.
2. At the LIB-44xx CLI prompt, press the **Enter** key.
3. Enter the CLI command “**config default keep_ip**” and press the **Enter** key. For example:

```
>config default keep_ip
>
```
4. Try reconnecting in IE.
5. If necessary, try another supported web browser.

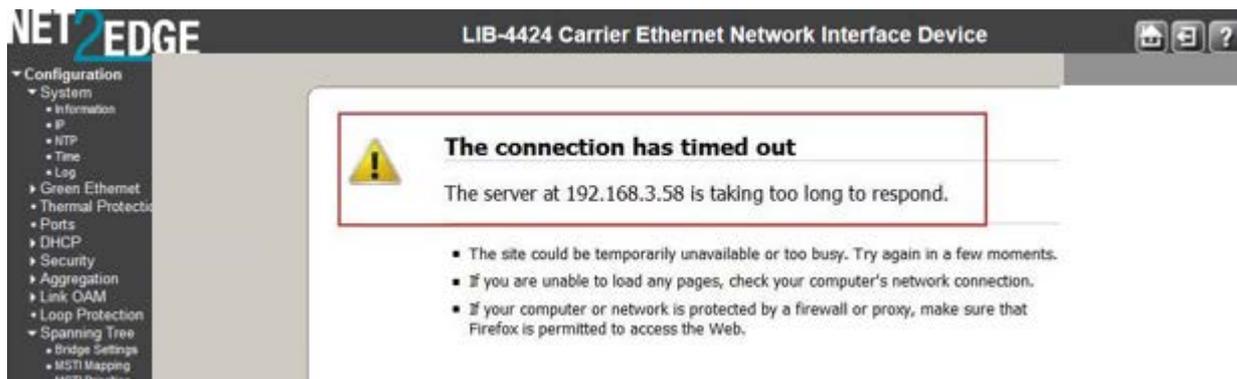
Problem: The LIB-44xx screen locks up or displays incorrectly in Internet Explorer.

Description: Internet Explorer 9 has some glitches in its operations. Some web pages do not render correctly with IE because the web pages were designed for earlier functions.

Resolution:

1. Enable compatibility view in Internet Explorer (at IE > **Tools** > **Compatibility View**).
2. Continue operation.

Message: **The Connection has timed out**



Meaning: The LIB-44xx switch cannot communicate with the PC after enabling Flow Control from the **Configuration** > **Port** > **Configuration** menu path. Increasing the CONSOLE Port Max Frame size to the tagged frame length fixes the problem.

Recovery:

1. Change the MGMT port to 1522 as the default frame size.

Appendix A - PCB Configurations

The PCB has configurable components that should only be altered by or at the direction of a Tech Support specialist.

See the “[Service](#)” [section](#) on page [512](#).

Appendix B - Licenses

This appendix provides LIB-44xx license information. At the **Monitor > System > Information** menu path you can click the **Acknowledgments > Details** link to display the current set of source code from the various Open-Source components.

Licenses

Name CPU-load
Description SVG graph
License type BSD

Copyright (C) 2004-2005 T. Lechat <dev@lechat.org>, Manuel Kasper <mk@neon1.net> and Jonathan Watt <jwatt@jwatt.org>.
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Name Dropbear
Description SSH Server
License type MIT, BSD, OpenSSL

Dropbear contains a number of components from different sources, hence there are a few licenses and authors involved. All licenses are fairly non-restrictive.

The majority of code is written by Matt Johnston, under the license below:

Portions of the client-mode work are (c) 2004 Mihnea Stoenescu, under the same license:

Copyright (c) 2002-2006 Matt Johnston
Portions copyright (c) 2004 Mihnea Stoenescu
All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

=====

LibTomCrypt and LibTomMath are written by Tom St Denis, and are Public Domain.

=====

sshpty.c is taken from OpenSSH 3.5p1,
Copyright (c) 1995 Tatu Ylonen <ylo@cs.hut.fi>, Espoo, Finland
All rights reserved

"As far as I am concerned, the code I have written for this software can be used freely for any purpose. Any derived versions of this software must be clearly marked as such, and if the derived work is incompatible with the protocol description in the RFC file, it must be called by a name other than "ssh" or "Secure Shell". "

=====

loginrec.c
loginrec.h
atomicio.h
atomicio.c
and strlcat() (included in util.c) are from OpenSSH 3.6.1p2, and are licensed under the 2 point BSD license.

loginrec is written primarily by Andre Lucas, atomicio.c by Theo de Raadt.

strlcat() is (c) Todd C. Miller

=====

Import code in keyimport.c is modified from PuTTY's import.c, licensed as follows:

PuTTY is copyright 1997-2003 Simon Tatham.

Portions copyright Robert de Bath, Joris van Rantwijk, Delian Delchev, Andreas Schultz, Jeroen Massar, Wez Furlong, Nicolas Barry, Justin Bradford, and CORE SDI S.A.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Name Host AP
Description EAPOL Authenticator and RADIUS authentication server
License type BSD

Copyright (c) 2002-2012, Jouni Malinen <j@w1.fi> and contributors
All Rights Reserved.

This program is licensed under the BSD license (the one with advertisement clause removed).

License

This software may be distributed, used, and modified under the terms of BSD license:

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name(s) of the above-listed copyright holder(s) nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Name ISC DHCP
Description DHCP Relay
License type ISC

Copyright (c) 2004-2008 by Internet Systems Consortium, ("ISC")
Copyright (c) 1995-2003 by Internet Software Consortium

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND ISC DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL ISC BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Name MD5
Description MD5 hash implementation
License type BSD

Copyright (c) 2003-2005, Jouni Malinen <j@w1.fi>

This program is free software; you can redistribute it and/or modify

it under the terms of the GNU General Public License version 2 as published by the Free Software Foundation.

Alternatively, this software may be distributed under the terms of BSD license.

 Name MooTools
 Description JavaScript Framework
 License type MIT

The MIT License

Copyright (c) 2006-2009 Valerio Proietti, <<http://mad4milk.net/>>

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

 Name NET-SNMP
 Description SNMP Agent
 License type NET-SNMP (BSD-Style)

Various copyrights apply to this package, listed in 5 separate parts below: Please make sure that you read all the parts. Up until 2001, the project was based at UC Davis, and the first part covers all code written during this time. From 2001 onwards, the project has been based at SourceForge, and Networks Associates Technology, Inc. hold the copyright on behalf of the wider Net-SNMP community, covering all derivative work done since then. An additional copyright section has been added as Part 3 below also under a BSD license for the work contributed by Cambridge Broadband Ltd. to the project since 2001. An additional copyright section has been added as Part 4 below also under a BSD license for the work contributed by Sun Microsystems, to the project since 2003.

Code has been contributed to this project by many people over the years it has been in development, and a full list of contributors can be found in the README file under the THANKS section.

---- Part 1: CMU/UCD copyright notice: (BSD like) ----

Copyright 1989, 1991, 1992 by Carnegie Mellon University

Derivative Work - 1996, 1998-2000

Copyright 1996, 1998-2000 The Regents of the University of California

All Rights Reserved

Permission to use, copy, modify and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appears in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of CMU and The Regents of the University of California not be used in advertising or publicity pertaining to distribution of the software without specific written permission.

CMU AND THE REGENTS OF THE UNIVERSITY OF CALIFORNIA DISCLAIM ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL CMU OR THE REGENTS OF THE UNIVERSITY OF CALIFORNIA BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM THE LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

---- Part 2: Networks Associates Technology, Inc. copyright notice (BSD) ----

Copyright (c) 2001-2003, Networks Associates Technology, Inc.

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- * Neither the name of the Networks Associates Technology, Inc. nor the names of its contributors may be used to endorse or promote

products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 3: Cambridge Broadband Ltd. copyright notice (BSD) ----

Portions of this code are copyright (c) 2001-2003, Cambridge Broadband Ltd. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- * The name of Cambridge Broadband Ltd. may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDER "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 4: Sun Microsystems, copyright notice (BSD) ----

Copyright © 2003 Sun Microsystems, , 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

Use is subject to license terms below:

This distribution may include materials developed by third parties.

Sun, Sun Microsystems, the Sun logo and Solaris are trademarks or registered trademarks of Sun Microsystems, in the U.S. and other countries.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- * Neither the name of the Sun Microsystems, nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 5: Sparta, Inc. copyright notice (BSD) -----

Copyright (c) 2003-2004, Sparta, Inc.
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- * Neither the name of Sparta, Inc. nor the names of its contributors may

be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Name NET-SNMP RMON
Description NET-SNMP RMON utilities
License type Alex Rozin, Optical Access

Copyright (C) 2001 Alex Rozin, Optical Access

All Rights Reserved

Permission to use, copy, modify and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation.

ALEX ROZIN DISCLAIM ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS, IN NO EVENT SHALL ALEX ROZIN BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Name NTP - Network Time Protocol
Description NTP Protocol
License type NTP

Copyright (c) David L. Mills 1992-2009

Permission to use, copy, modify, and distribute this software and its documentation for any purpose with or without fee is hereby granted, provided that the above copyright notice appears in all copies and that both the copyright notice and this permission notice appear in supporting documentation, and that the name University of Delaware not be used in advertising or publicity

pertaining to distribution of the software without specific, written prior permission. The University of Delaware makes no representations about the suitability this software for any purpose. It is provided "as is" without express or implied warranty.

Copyright (c) 1990, 1993

The Regents of the University of California. All rights reserved.

This code is derived from software contributed to Berkeley by Chris Torek.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement:

This product includes software developed by the University of California, Berkeley and its contributors.

4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Copyright (c) 1987, 1989 Regents of the University of California.
All rights reserved.

This code is derived from software contributed to Berkeley by Arthur David Olson of the National Cancer Institute.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

3. All advertising materials mentioning features or use of this software must display the following acknowledgement:

This product includes software developed by the University of California, Berkeley and its contributors.

4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Copyright (c) 1983, 1993

The Regents of the University of California. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

3. All advertising materials mentioning features or use of this software must display the following acknowledgement:

This product includes software developed by the University of California, Berkeley and its contributors.

4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY

OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Amanda, The Advanced Maryland Automatic Network Disk Archiver
Copyright (c) 1991-1998 University of Maryland at College Park
All Rights Reserved.

Permission to use, copy, modify, distribute, and sell this software and its documentation for any purpose is hereby granted without fee, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of U.M. not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. U.M. makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

U.M. DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS, IN NO EVENT SHALL U.M. BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Author: James da Silva, Systems Design and Analysis Group
Computer Science Department
University of Maryland at College Park

Copyright (C) 1991-2, RSA Data Security, Created 1991. All rights reserved.

License to copy and use this software is granted provided that it is identified as the "RSA Data Security, MD5 Message-Digest Algorithm" in all material mentioning or referencing this software or this function.

License is also granted to make and use derivative works provided that such works are identified as "derived from the RSA Data Security, MD5 Message-Digest Algorithm" in all material mentioning or referencing the derived work.

RSA Data Security, makes no representations concerning either the merchantability of this software or the suitability of this software for any particular purpose. It is provided "as is" without express or implied warranty of any kind.

These notices must be retained in any copies of any part of this documentation and/or software.

Name OpenSSL

Description Toolkit implementing SSL v2/v3 and TLS protocols
License type OpenSSL

Copyright (c) 2001 The OpenSSL Project. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment:
"This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (<http://www.openssl.org/>)"
4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact openssl-core@openssl.org.
5. Products derived from this software may not be called "OpenSSL" nor may "OpenSSL" appear in their names without prior written permission of the OpenSSL Project.
6. Redistributions of any form whatsoever must retain the following acknowledgment:
"This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>)"

THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT "AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com). This product includes software written by Tim

Hudson (tjh@cryptsoft.com).

Name PTPd

Description Precision Time protocol (PTP)

License type MIT-Style

Copyright (c) 2005-2008 Kendall Correll, Aidan Williams

Permission to use, copy, modify, and/or distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND THE AUTHOR DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Name WPA Supplicant

Description EAP Peer state machines for MAC-based Authentication

License type BSD

Copyright (c) 2003-2012, Jouni Malinen <j@w1.fi> and contributors
All Rights Reserved.

This program is licensed under the BSD license (the one with advertisement clause removed).

License

This software may be distributed, used, and modified under the terms of BSD license:

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name(s) of the above-listed copyright holder(s) nor the names of its contributors may be used to endorse or promote products

derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

 Name avltree
 Description Self-balancing binary search tree
 License type MIT

Copyright (c) 2011 Bijal Thanawala

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

 Name eCos RTOS
 Description Real-time OS for embedded applications
 License type Modified GPL

This file is part of eCos, the Embedded Configurable Operating System. Copyright (C) 1998, 1999, 2000, 2001, 2002, 2003, 2004, 2005, 2006, 2007, 2008, 2009 Free Software Foundation,

eCos is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 or (at your option) any later

version.

eCos is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with eCos; if not, write to the Free Software Foundation, , 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301, USA.

As a special exception, if other files instantiate templates or use macros or inline functions from this file, or you compile this file and link it with other works to produce a work based on this file, this file does not by itself cause the resulting work to be covered by the GNU General Public License. However the source code for this file must still be made available in accordance with section (3) of the GNU General Public License v2.

This exception does not invalidate any other reasons why a work based on this file might be covered by the GNU General Public License.

Appendix C - Application Notes and Use Cases

This appendix provides LIB-44xx use cases and related application notes.

LIB-44xx Applications Support

The LIB-44xx provides a rich set of Carrier Ethernet services, Ethernet switching, and Ethernet transport features. Advanced TCAM-based QoS processing enables delivery of differentiated services with per-service SLA guarantees.

The LIB-44xx can be used to implement the access and first-level aggregation functions in CPE (Customer Premise equipment) and PE (Provider edge) equipment placed in a network. In these applications, the LIB-44xx forms the demarcation between the customer network and the provider network, providing the full set of UNI features in a cost-effective manner:

- Map customer frame formats into Provider frame formats.
- Classify frames and map to appropriate QoS profiles.
- Apply Provider Bridge (PB) or Provider Backbone Bridge (PBB) encapsulations.
- Meter the customer traffic and ensure that the customer SLA (Service Level Agreement) is met.
- Police using MEF-defined Dual Leaky Bucket algorithm.
- Mark frames as Committed (Green) or Discard Eligible (Yellow).
- Provide correct QoS treatment (traffic management).
- Provide traffic statistics per Customer in a manner consistent with the SLA.
- Enable end-to-end Service OAM (SOAM) by the customer, if allowed.
- Implement the Service as defined by the SLA.
- E-LINE for point-to-point or backhaul services.
- E-LAN for multi-point/bridged services.
- Enable management and protection schemes as required by the Provider.
- Link Aggregation or other port protection schemes if used for access.
- OAM at the Operator and Service Provider levels for remote management, fault diagnosis, and protection switching.
- Support network timing and synchronization requirements as required.
- Provide Sync-E (Synchronous Ethernet) and IEEE 1588 PTP functionality.

Not Intended for Use in Life Support Products: LIB-44xx products are not intended for use in life support products, systems, or environments where failure of an LIB-44xx product could reasonably be expected to result in death or personal injury. Anyone using an LIB-44xx product in such an application without express written consent of an officer of Net2Edge Limited, does so at their own risk, and agrees to fully indemnify Net2Edge Limited, for any damages that may result from such use or sale.

Appendix D: Service, Warranty & Compliance Information

Service

Direct Contact Numbers:

Tel: +44-345-0130030, xtn 6810

Email: support@Net2Edge.com

Warranty

This warranty is your only remedy. No other warranties, such as fitness for a particular purpose, are expressed or implied. Net2Edge is not liable for any special, indirect, incidental or consequential damages or losses, including loss of data, arising from any cause or theory. Authorized resellers are not authorized to extend any different warranty on Net2Edge's behalf.

1 Year Limited Warranty

Net2Edge labelled LIB-44xx series products are warranted to be free from defects in material and workmanship for a period of 1 year beyond the Net2Edge shipment date. This warranty covers the original user only and is not transferable.

What the Warranty Does Not Cover

This warranty does not cover damage from accident, acts of God, neglect, contamination, misuse or abnormal conditions of operation or handling, including over-voltage failures caused by use outside the product's specified rating, or normal wear and tear of mechanical components.

Establishing Original Ownership

To establish original ownership and provide date of purchase, please complete and return the registration card accompanying the product or register the product on-line on our product registration page.

Net2Edge will at its option:

Repair the defective product to functional specifications at no charge

Replace the product with an equivalent functional product

Refund the purchase price of a defective product

Who to Contact for
Returns

Return Authorization

To return a defective product for warranty coverage, contact Net2Edge's technical support department for a return authorization number.

Return Instructions

Send the defective product postage and insurance prepaid to the following address:

Net2Edge
Kulite House,
Stroudley Road,
Basingstoke
RG24 8UG, UK.
Tel: +44 345 0130030
Attn: RETURNS DEPT: CRA/RMA # _____

Failure to properly protect the product during shipping may void this warranty. The return authorization number must be written on the outside of the carton to ensure its acceptance. We cannot accept delivery of any equipment that is sent to us without a CRA or RMA number.

CRA's are valid for 60 days from the date of issuance. An invoice will be generated for payment on any unit(s) not returned within 60 days.

Upon completion of a demo/ evaluation test period, units must be returned or purchased within 30 days. An invoice will be generated for payment on any unit(s) not returned within 30 days after the demo/ evaluation period has expired.

The customer must pay for the non-compliant product(s) return transportation costs to Net2Edge for evaluation of said product(s) for repair or replacement. Net2Edge will pay for the shipping of the repaired or replaced in-warranty product(s) back to the customer (any and all customs charges, tariffs, or/and taxes are the customer's responsibility).

Before making any non-warranty repair, Net2Edge requires a \$200.00 charge plus actual shipping costs to and from the customer. If the repair is greater than \$200.00, an estimate is issued to the customer for authorization of repair. If no authorization is obtained, or the product is deemed 'not repairable', Net2Edge will retain the \$200.00 service charge and return the product to the customer not repaired. Non-warranted products that are repaired by Net2Edge for a fee will carry a 180-day limited warranty. All warranty claims are subject to the restrictions and conventions set forth by this document.

Net2Edge reserves the right to charge for all testing and shipping incurred, if after testing, a return is classified as "No Problem Found."

THIS WARRANTY IS YOUR ONLY REMEDY. NO OTHER WARRANTIES, SUCH AS FITNESS FOR A PARTICULAR PURPOSE, ARE EXPRESSED OR IMPLIED. NET2EDGE IS NOT LIABLE FOR ANY SPECIAL, INDIRECT, INCIDENTAL OR CONSEQUENTIAL DAMAGES OR LOSSES, INCLUDING LOSS OF DATA, ARISING FROM ANY CAUSE OR THEORY. AUTHORIZED RESELLERS ARE NOT AUTHORIZED TO EXTEND ANY DIFFERENT WARRANTY ON NET2EDGE'S BEHALF.

Customer Pays Non-Compliant Return Costs The customer must pay the non-compliant product(s) return transportation cost to Net2Edge for evaluation of said product(s) for repair or replacement. Net2Edge will pay for shipping the repaired or replaced in-warranty product(s) back to the customer (any and all customs charges, tariffs, or/and taxes are the customer's responsibility).

Non-Warranty Repair Costs Before making any non-warranty repair, Net2Edge requires a \$200 charge, plus actual shipping costs to and from the customer. If the repair is greater than \$200, an estimate is issued to the customer for authorization before making the repair. If no authorization is obtained, or the product is deemed not repairable, Net2Edge will retain the \$200 service charge and return the product to the customer not repaired.

Repaired Non-Warranty Products Non-warranted products repaired by Net2Edge for a fee will carry a 180-day limited warranty. All warranty claims are subject to the restrictions and conventions set forth by this document. Net2Edge reserves the right to charge for all testing and shipping incurred, if after testing, a return is classified as "No Problem Found."

Compliance Information

Standards: CISPR22/EN55022 Class A, CE Mark

FCC Regulations:

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

This Class A digital apparatus complies with Canadian ICES-003.

Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada.

CE Marking

This is a Class A product. In a domestic environment, this product could cause radio interference; as a result, the customer may be required to take adequate preventative measures.

UL Recognized

Tested and recognized by the Underwriters Laboratories,

European Regulations

WARNING:

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

Achtung !

Dieses ist ein Gerät der Funkstörgrenzwertklasse A. In Wohnbereichen können bei Betrieb dieses Gerätes Rundfunkstörungen auftreten. In diesem Fall ist der Benutzer für Gegenmaßnahmen verantwortlich.

Attention !

Ceci est un produit de Classe A. Dans un environnement domestique, ce produit risque de créer des interférences radioélectriques, il appartiendra alors à l'utilisateur de prendre les mesures spécifiques appropriées.



In accordance with European Union Directive 2002/96/EC of the European Parliament and of the Council of 27 January 2003, Net2Edge will accept post usage returns of this product for proper disposal. The contact information for this activity can be found in the 'Contact Us' portion of this document.



CAUTION: RJ connectors are NOT INTENDED FOR CONNECTION TO THE PUBLIC TELEPHONE NETWORK. Failure to observe this caution could result in damage to the public telephone network.

Der Anschluss dieses Gerätes an ein öffentliches Telekommunikationsnetz in den EG-Mitgliedstaaten verstösst gegen die jeweiligen einzelstaatlichen Gesetze zur Anwendung der Richtlinie 91/263/EWG zur Angleichung der Rechtsvorschriften der Mitgliedstaaten über Telekommunikationsendeinrichtungen einschliesslich der gegenseitigen Anerkennung ihrer Konformität.

Declaration of Conformity

Declaration of Conformity

Transition Networks, Inc.

Manufacture's Name

10900 Red Circle Drive, Minnetonka, Minnesota 55343 U.S.A.

Manufacture's Address

Declares that the products:

S4140, S4212, S4224

Conforms to the following Product Regulations:

EN 55022:2010/AC:2011, EN 55024:2010, IEC /EN 60950-1

Directive 2004/108/EC

FCC Part 15 Subpart B

AS/NZS CISPR 22: 2009 + A1:2010, ICES-003, Issue 5:2012, VCCI V-3/2013.04

I, the undersigned, hereby declare that the equipment specified above conforms to the above Directive(s) and Standards(s).

Minnetonka, Minnesota

Place

Dec 23, 2013

Date

Stephen Anderson

Signature

Stephen Anderson

Full Name

Vice President of Engineering

Position

20141B

Electrical Safety Warnings

Electrical Safety

IMPORTANT: This equipment must be installed in accordance with safety precautions.

Elektrische Sicherheit

WICHTIG: Für die Installation dieses Gerätes ist die Einhaltung von Sicherheitsvorkehrungen erforderlich.

Elektrisk sikkerhed

VIGTIGT: Dette udstyr skal installeres i overensstemmelse med sikkerhedsadvarslerne.

Elektrische veiligheid

BELANGRIJK: Dit apparaat moet in overeenstemming met de veiligheidsvoorschriften worden geïnstalleerd.

Sécurité électrique

IMPORTANT : Cet équipement doit être utilisé conformément aux instructions de sécurité.

Sähköturvallisuus

TÄRKEÄÄ : Tämä laite on asennettava turvaohjeiden mukaisesti.

Sicurezza elettrica

IMPORTANTE: questa apparecchiatura deve essere installata rispettando le norme di sicurezza.

Elektrisk sikkerhet

VIKTIG: Dette utstyret skal installeres i samsvar med sikkerhetsregler.

Segurança eléctrica

IMPORTANTE: Este equipamento tem que ser instalado segundo as medidas de precaução de segurança.

Seguridad eléctrica

IMPORTANTE: La instalación de este equipo deberá llevarse a cabo cumpliendo con las precauciones de seguridad.

Elsäkerhet

OBS! Alla nödvändiga försiktighetsåtgärder måste vidtas när denna utrustning används

Safety Instructions for Rack Mount Installations

The instructions below (or similar) are intended for LIB-44xx rackmount installation environments:

1. Elevated Operating Ambient: if installed in a closed or multi-unit rack assembly, the operating ambient temperature of the rack environment may exceed room ambient. Install the equipment in an environment compatible with the maximum ambient temperature (T_{ma}) specified.
2. Reduced Air Flow: install the equipment in a rack so that the amount of air flow required for safe operation is not compromised.
3. Mechanical Loading: Mount the equipment in the rack so that a hazardous condition does not occur due to uneven mechanical loading (weight distribution/rack balance).
4. Circuit Overloading: give consideration to the connection of the equipment to the supply circuit and the effect that overloading of the circuits might have on overcurrent protection and supply wiring. Consider all equipment nameplate ratings when addressing this concern.
5. Reliable Earthing: maintain reliable earthing of rack-mounted equipment; pay particular attention to supply connections other than direct connections to the branch circuit (e.g., use of power strips).

Appendix E: SNMP MIBs and Traps

This appendix describes the LIB-44xx SNMP MIBs and Traps.

Supported MIBs

The LIB-44xx supports public (standard) and private Management Information Bases (MIBs). The LIB-44xx public MIBs are listed below:

Table 9: Public MIBs

MIB	Note
RFC 1213 MIB II	MIB for Network Management of TCP/IP-based internets: MIB-II . Defines the second version of the MIB-II for use with network management protocols in TCP/IP-based internets.
IP-MIB	Request for Comments: 4293 . PROPOSED STANDARD; Errata Exist. One primary purpose of this revision of the IP MIB is to create a single set of objects to describe and manage IP modules in an IP version independent manner.
RFC 4188 Bridge MIB	Definitions of Managed Objects for Bridges. This RFC defines a portion of the MIB for use with network management protocols in TCP/IP-based internets. In particular, it defines objects for managing MAC bridges based on the IEEE 802.1D-1998 standard between LAN segments. Provisions are made for the support of transparent bridging. Provisions are also made so that these objects apply to bridges connected by subnetworks other than LAN segments.
RFC 2674 VLAN MIB	Definitions of Managed Objects for Bridges with Traffic Classes, Multicast Filtering and Virtual LAN Extensions. RFC 2674 defines two MIB modules for managing the new capabilities of MAC bridges defined by the IEEE 802.1D-1998 MAC Bridges and the IEEE 802.1Q-1998 Virtual LAN (VLAN) standards for bridging between LAN)segments. One MIB module defines objects for managing the 'Traffic Classes' and 'Enhanced Multicast Filtering' components of IEEE 802.1D-1998. The other MIB module defines objects for managing IEEE 802.1Q VLANs.
RFC 4878 Link OAM MIB	Definitions and Managed Objects for OAM Functions on Ethernet-Like Interfaces. RFC 4878 defines objects for controlling link OAM functions and for providing results and status of the OAM functions to management entities.
IEEE 802.1AB (LLDP MIB)	802.1AB-2009 - IEEE Standard for LANs/MANs-- Station and Media Access Control Connectivity Discovery defines a protocol and a set of managed objects that can be used for discovering the physical topology from adjacent stations in IEEE 802 LANs.
IEEE 802.1 (MSTP MIB)	IEEE 802.1™: BRIDGING & MANAGEMENT; for the full set of IEEE Standards for Local and metropolitan area networks see http://standards.ieee.org/about/get/802/802.1.html .
IEEE 802.1X (PAE MIB)	802.1X - Port Based Network Access Control; defines the changes necessary to the operation of a MAC Bridge in order to provide Port based network access control capability.

IEEE 802.30ad (LACP MIB)	Amendment to Carrier Sense Multiple Access With Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications-Aggregation of Multiple Link Segments.
RFC 2819 RMON	(Group 1, 2, 3 & 9.) Remote Network Monitoring MIB; defines a portion of the MIB for use with network management protocols in TCP/IP-based internets. In particular, it defines objects for managing remote network monitoring devices.
RFC 2613 SMON MIB	Remote Network Monitoring MIB Extensions for Switched Networks Version 1.0; defines a portion of the MIB for use with network management protocols in TCP/IP-based internets. In particular, it defines objects for managing remote network monitoring devices in switched networks environments.
RFC 2863 Interface Group MIB	The Interfaces Group MIB defines a portion of the Management Information Base (MIB) for use with network management protocols in the Internet community. In particular, it describes managed objects used for managing Network Interfaces.
RFC 3635 Ethernet-like MIB	Definitions of Managed Objects for the Ethernet-like Interface Types; defines objects for managing Ethernet-like interfaces; it updates RFC 2665 by including management information useful for the management of 10 Gigabit per second (Gb/s) Ethernet interfaces.
RFC 3636 802.3 MAU MIB	Definitions of Managed Objects for IEEE 802.3 Medium Attachment Units (MAUs); defines objects for managing IEEE 802.3 MAUs; extends earlier specifications by including management information for the management of 10 gigabit per second (Gb/s) MAUs.
RFC 4133 Entity MIB version 3	Entity MIB (Version 3) ; defines a portion of the MIB for use with network management protocols in the Internet community. It describes managed objects used for managing multiple logical and physical entities managed by a single SNMP agent. RFC 4133 specifies version 3 of the Entity MIB, which obsoletes version 2 (RFC 2737).
RFC 4668 RADIUS Auth. Client MIB	RADIUS Authentication Client MIB for IPv6; defines a set of extensions that instrument RADIUS authentication client functions. These extensions represent a portion of the MIB for use with network management protocols in the Internet community. Using these extensions, IP-based management stations can manage RADIUS authentication clients. RFC 4668 obsoletes RFC 2618 by deprecating the MIB table containing IPv4-only address formats and defining a new table to add support for version-neutral IP address formats.
RFC 3411 SNMP Management Frameworks	An Architecture for Describing SNMP Management Frameworks ; describes an architecture for describing Simple Network Management Protocol (SNMP) Management Frameworks. The architecture is designed to be modular to allow the evolution of the SNMP protocol standards over time.

SNMPv3 MPD	RFC 2572 (Message Processing and Dispatching for the Simple Network Management Protocol (SNMP) describes the Message Processing and Dispatching for SNMP messages within the SNMP architecture [RFC2571]. It defines the procedures for dispatching potentially multiple versions of SNMP messages to the proper SNMP Message Processing Models, and for dispatching PDUs to SNMP applications. This document also describes one Message Processing Model - the SNMPv3 Message Processing Model.
RFC 3414 USM SNMPv3	User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3); defines the Elements of Procedure for providing SNMP message level security. RFC 3414 also includes a MIB for remotely monitoring/managing the configuration parameters for this Security Model.
RFC 3415 VACM SNMPv3	VACM (View-based Access Control Model) for the Simple Network Management Protocol (SNMP); defines the Elements of Procedure for controlling access to management information. This document also includes a MIB for remotely managing the configuration parameters for the View-based Access Control Model. RFC 3415 obsoletes RFC 2575.
IEEE 802.1AG MIB	The CFM (Connectivity Fault Management) standard specifies protocols, procedures, and managed objects to support transport fault management. These allow discovery and verification of the path, through bridges and LANs, taken for frames addressed to and from specified network users, detection, and isolation of a connectivity fault to a specific bridge or LAN. This standard will provide capabilities for detecting, verifying and isolating connectivity failures in networks.
MEF 31 (SOAM FM)	Service OAM Fault Management Definition of Managed Objects; specifies the Fault Management (FM) MIB necessary to implement the Service Operations, Administration, and Maintenance (OAM) that satisfies the Service OAM requirements and framework specified by MEF 17, the Service OAM Fault Management requirements as specified by SOAM-FM, and the Service OAM management objects as specified by MEF 7.1 which are applicable to FM functions. Two non-MEF documents serve as the baseline documents for this work: ITU-T Y.1731 and IEEE 802.1ag.
MEF SOAM PM (draft)	MEF 36 - Service OAM SNMP MIB for Performance Monitoring; specifies the Performance Monitoring (PM) MIB necessary to manage Service Operations, Administration, and Maintenance (OAM) implementations that satisfy the Service OAM requirements and framework specified by MEF 17, the Service OAM Performance Monitoring requirements as specified by SOAM-PM, and the Service OAM management objects as specified by MEF 7.1 which are applicable to PM functions. Two non-MEF documents serve as the baseline documents for this work: ITU-T Y.1731 and IEEE 802.1ag.

SNMP Traps List

The table below lists and describes the Trap MIB variables.

SNMP Traps List

Table 12: Traps List

No	Trap MIB Variable	Binding	OID	Cause
1	SNMPv2-MIB:coldStart	NULL	1.3.6.1.6.3.1.1.5.1	When the device undergoes a reboot.
2	SNMPv2-MIB:warmStart	NULL	1.3.6.1.6.3.1.1.5.2	<Not implemented>
3	SNMPv2-MIB:linkDown	1: ifIndex 2: ifAdminStatus 3: ifOperStatus	1.3.6.1.6.3.1.1.5.3	When a port's link goes down due to adminstate change or due to physical layer connection.
4	SNMPv2-MIB:linkUp	1: ifIndex 2: ifAdminStatus 3: ifOperStatus	1.3.6.1.6.3.1.1.5.4	When a port's link goes up.
5	SNMPv2-MIB:authenticationFailure	NULL	1.3.6.1.6.3.1.1.5.5	When SNMP community string sent in a request doesn't match the configured community string.
6	LLDP-MIB:lldpRemTablesChange	1: lldpStatsRemTablesInserts 2: lldpStatsRemTablesDeletes 3: lldpStatsRemTablesDrops 4: lldpStatsRemTablesAgeouts	1.0.8802.1.1.2.0.0.1	When the topology of connected remote devices changes.
7	BRIDGE-MIB:newRoot	NULL	1.3.6.1.2.1.17.0.1	When the agent becomes the new Root of Spanning tree.
8	BRIDGE-MIB:topologyChange	NULL	1.3.6.1.2.1.17.0.2	When the status of configured ports in a Bridge changes from Learning to Forwarding state, or from Forwarding to Blocking state.
9	LLDP-EXT-MED-MIB:lldpXMedTopologyChangeDetected	1: lldpRemChassisIdSubtype 2: lldpRemChassisId 3: lldpXMedRemDeviceClass	1.0.8802.1.1.2.1.5.4795.0.1	When a new remote device is connected or disconnect from the local device.
10	ENTITY-MIB:entConfigChange	NULL	1.3.6.1.2.1.47.2.0.1	When the device entity changes.
11	TN-MGMT-MIB:tnDMIRxIntrusionEvt	1: ifIndex 2: tnDMIRxPwrLvlPreset 3: tnDMIRxPowerLevel	1.3.6.1.4.1.868.2.5.3.0.1	When the tnDMIRxPowerLevel falls below the tnDMIRxPwrLvlPreset, indicating an intrusion on the fiber.
12	TN-MGMT-MIB:tnDMIRxPowerEvt	1: ifIndex 2: tnDMIRxPowerAlarm 3: tnDMIRxPowerLevel	1.3.6.1.4.1.868.2.5.3.0.2	When there is a warning or alarm on Rx Power.
13	TN-MGMT-MIB:tnDMITxPowerEvt	1: ifIndex 2: tnDMITxPowerAlarm 3: tnDMITxPowerLevel	1.3.6.1.4.1.868.2.5.3.0.3	When there is a warning or alarm on Tx Power.
14	TN-MGMT-MIB:tnDMITxBiasEvt	1: ifIndex 2: tnDMITxBiasAlarm 3: tnDMITxBiasCurrent	1.3.6.1.4.1.868.2.5.3.0.4	When there is a warning or alarm on Tx Bias current.
15	TN-MGMT-MIB:tnDMITemperatureEvt	1: ifIndex 2: tnDMITempAlarm 3: tnDMITemperature	1.3.6.1.4.1.868.2.5.3.0.5	When there is a warning or alarm on DMI temperature.
16	TN-MGMT-MIB:tnflLimitDynMACEvt	1: ifIndex 2: tnflLimitDynMACMaxCount 3: tnflLimitDynMACState	1.3.6.1.4.1.868.2.5.3.0.8	When a port which has Limit control on dynamic MAC is enabled and the limit is reached.
17	TN-LOOP-PROTECT-MIB:tnLoopProtectLoopDetectedNotification	1: ifIndex 2: tnLoopProtectPortLoopCount 3: tnLoopProtectPortAction	1.3.6.1.4.1.868.2.5.22.0.1	When a loop is detected in a port.
18	TN-ELPS-MIB:tnElpsAlarm	1: tnElpsWFlowState 2: tnElpsPFlowState 3: tnElpsArchitectureMismatch 4: tnElpsAPSONWorking 5: tnElpsSwitchingIncomplete	1.3.6.1.4.1.868.2.5.109.2.0.1	When the configuration of EPS is modified.
19	NOTIFICATION-TYPE:tnThermalProtectionPortStatusChangedNotification	1: ifIndex 2: tnThermalProtectionPriorityTemperature 3: tnThermalProtectionIfStatusTemperature 4: tnThermalProtectionIfStatusCode	1.3.6.1.4.1.868.2.5.32.0.1	A notification generated by the local device sensing a change in the thermal protection port status. The change indicates the current temperature of a port becomes higher (or lower) than its priority temperature.

Private MIB OID Assignments

OID	MIB module Name
tnProducts.3	tnMgmtMIB
tnProducts.4	tnVlanQoSManagementMIB
tnProducts.6	tnEntitySensorMIB
tnProducts.7	tnLOAMExtMIB
tnProducts.9	tnQosExtMIB
tnProducts.10	tnDevSysIpMgmtMIB
tnProducts.19	tnSysCfgChangeMIB
tnProducts.20	tnPowerSupply
tnProducts.21	tnFraMIB
tnProducts.22	tnLoopProtectMIB
tnProducts.25	tnMirroringMIB
tnProducts.26	tnPrivateVlanMIB
tnProducts.31	tnIPSourceGuardMIB
tnProducts.32	tnThermalProtectionMIB
tnProducts.33	tnDhcpMIB
tnProducts.50	tnMVRMIB
tnProducts.60	tnCesMIB
tnProducts.61	tnCesRoutingMIB
tnProducts.105	tnEthSoamMIB
tnProducts.106	tnEvcMIB
tnProducts.107	tnEvcIdentityMIB
tnProducts.108	tnEtherSAT
tnProducts.109	tnProtectionMIB
tnProducts.110	tnProvMIB
tnProducts.111	tnXstpMib
tnProducts.114	tnIcmpSnoopingMib
tnProducts.115	tnMldSnoopingMib
tnProducts.119	tnSFlowMIB
tnProducts.120	tnMRP
tnProducts.121	tnStaticIpRouting
tnProducts.122	tnSynceMIB
tnProducts.123	tnPtpMIB
tnProducts.125	tnNASMIB
tnProducts.126	tnZeroTouchProvisionMIB
tnProducts.130	tnEthSatLoopbackMIB
tnProducts.137	tnExtLldpMIB
tnProducts.138	tnAggCfgMIB

tnProducts.140	tnLinkOamMIB
tnProducts.141	tnPortMIB
tnProducts.142	tnMacMib

tnMgmtMIB subtree OID assignments

OID	MIB module Name
tnMgmtMIB.1.1	tnDevMgmt
tnMgmtMIB.1.2	tnInterfaceMgmt
tnMgmtMIB.1.3	tnInterfaceDiagMgmt
tnMgmtMIB.1.4	tnIfMACSecurityMgmt
tnMgmtMIB.1.5	tnIfQOSMgmt

tnDevMgmt subtree OID assignments

OID	MIB module Name
tnDevMgmt.1	tnDevSysMgmt
tnDevMgmt.2	tnDevSysLPT
tnDevMgmt.3	tnDevSysDyingGasp
tnDevMgmt.4	tnDevSysMACLearning
tnDevMgmt.5	tnAcIMgmt
tnDevMgmt.11	tnDevSysxNTP
tnDevMgmt.14	tnDevSysSnmpmgmt
tnDevMgmt.18	tnSyslogMIB
tnDevMgmt.19	tnDevSysUser
tnDevMgmt.21	tnSecurityAAAMIB
tnDevMgmt.22	tnARPIInspectionMIB
tnDevMgmt.30	tnDevSysUpgraderMIB
tnDevMgmt.36	tnDevAccessMgmtMIB
tnDevMgmt.37	tnDevVlanTranslation
tnDevMgmt.38	tnDevAggregation
tnDevMgmt.42	tnSecuritySwitchSSHMIB
tnDevMgmt.43	tnHttpsMib

Private MIBs

Table 10: Private MIBs

#	MIB File	Table	Version
1	-MGMT-MIB.smi	tnEthInterfaceTable	
2	-MGMT-MIB.smi	tnDevSysCfgTable	
3	-MGMT-MIB.smi	tnDevSysMacLearningTable	
4	-MGMT-MIB.smi	tnDMIIInfoTable	
5	-MGMT-MIB.smi	tnIfLimitDynMACLearningTable	
6	-MGMT-MIB.smi	tnIfTDRTestTable	
7	-MGMT-MIB.smi	tnIfTDRResultTable	
8	-PROV-MIB	tnProvTable	
9	-EVC-MIB	tnEvcPortTable	
10	-EVC-MIB	tnEvcBandwidthProfilesTable	
11	-VLAN-MGMT-MIB	tnSysManagmentVLANTable	
12	-VLAN-MGMT-MIB	tnSysVLANExtTable	
13	-VLAN-MGMT-MIB	tnIfVLANtagMgmt2Table	
14	P-BRIDGE-MIB	dot1dPortPriorityTable	
15	P-BRIDGE-MIB	dot1dUserPriorityRegenTable	
16	P-BRIDGE-MIB	dot1dTrafficClassTable	
17	P-BRIDGE-MIB	dot1dPortOutboundAccessPriorityTable	
18	ENTITY-MIB	entPhysicalTable	
19	-DEV-SYS-UPGRADER-MIB	tnFirmwareUpgradeTable	
20	-EVC-MIB	tnEvcTable	
21	-EVC-MIB	tnEvcEceTable	
22	-ETHSOAM-MIB	tnEthSoamMPTable	
23	IEEE8021-CFM-MIB	dot1agCfmMepTable	
24	MEF-SOAM-FM-MIB	mefSoamLckCfgTable	
25	MEF-SOAM-FM-MIB	mefSoamTestCfgTable	
26	MEF-SOAM-FM-MIB	mefSoamTestStatsTable	
27	MEF-SOAM-FM-MIB	mefSoamAisCfgTable	
28	MEF-SOAM-FM-MIB	mefSoamLmCfgTable	
29	MEF-SOAM-FM-MIB	mefSoamLbCfgTable	
30	MEF-SOAM-FM-MIB	mefSoamLbrMulticastTable	
31	-ETHSOAM-MIB	tnEthSoamLocalCfgTable	
32	-ETHSOAM-MIB	tnEthSoamStatusTable	
33	-ETHSOAM-MIB	tnEthSoamLossStateTable	
34	-ETHSOAM-MIB	tnEthSoamTSExtTable	
35	-ETHSOAM-MIB	tnEthSoamPeerCfgTable	
36	-ETHSOAM-MIB	tnEthSoamPeerStatusTable	
37	-ETHSOAM-MIB	tnEthSoamClientCfgTable	
38	-ETHSOAM-MIB	tnEthSoamLtmTable	
39	-ETHSOAM-MIB	tnEthSoamLtrTable	
40	IEEE8021-CFM-MIB	dot1agCfmLtrTable	
41	-ETHSOAM-MIB	tnEthSoamAisCfgTable	
42	-ETHSOAM-MIB	tnEthSoamDmCfgTable	
43	-ETHSOAM-MIB	tnEthSoamDmStateTable	
44	-QOS-EXT-MIB	tnQosExtPortPolicerTable	
45	-QOS-EXT-MIB	tnQosExtPortQueuePolicerTable	

46	-QOS-EXT-MIB	tnQosExtPortSchedulerTable	
47	-QOS-EXT-MIB	tnQosExtPortSchedulerWeightTable	
48	-QOS-EXT-MIB	tnQosExtPortShaperTable	
49	-QOS-EXT-MIB	tnQosExtPortQueueShaperTable	
50	-QOS-EXT-MIB	tnQosExtPortStormControlTable	
51	-DEV-SYS-IPMGMT-MIB	tnIpMgmtTable	
52	-DEV-SYS-IPMGMT-MIB	tnDnsServerTable	
53	-DEV-SYS-IPMGMT-MIB	tnIpextMgmtTable	
54	-SYS-LOG-MIB	tnSyslogMgmtTable	
55	-SYS-LOG-MIB	tnSyslogMessageTable	
56	-SYS-LOG-MIB	tnSyslogExtTable	
57	-MIRRORING-MIB	tnMirroringGroupTable	
58	-MIRRORING-MIB	tnMirroringIfTable	
59	-LOOP-PROTECT-MIB	tnLoopProtectBaseTable	
60	-LOOP-PROTECT-MIB	tnLoopProtectPortTable	
61	-DEV-SYS-XNTP-MIB	tnxNTPServerTable	
62	-PRIVATE-VLAN-MIB	tnPVlanMembershipTable	
63	-PRIVATE-VLAN-MIB	tnPVlanPortIsolationTable	
64	-EVC-MIB	tnEvcl2cpCfgTable	v1.4
65	IEEE8021-SPANNING-TREE-MIB	ieee8021SpanningTreeTable	v1.4
66	IEEE8021-SPANNING-TREE-MIB	ieee8021SpanningTreePortTable	v1.4
67	IEEE8021-MSTP-MIB	ieee8021MstpCistTable	v1.4
68	IEEE8021-MSTP-MIB	ieee8021MstpCistPortTable	v1.4
69	IEEE8021-MSTP-MIB	ieee8021MstpTable	v1.4
70	IEEE8021-MSTP-MIB	ieee8021MstpPortTable	v1.4
71	IEEE8021-MSTP-MIB	ieee8021MstpConfigIdTable	v1.4
72	IEEE8021-MSTP-MIB	ieee8021MstpVlanTable	v1.4
73	-XSTP-MIB	tnExtMstpCistTable	v1.4
74	-XSTP-MIB	tnExtMstpTable	v1.4
75	-XSTP-MIB	tnExtMstpCistPortTable	v1.4
76	-XSTP-MIB	tnExtMstpPortTable	v1.4
77	-XSTP-MIB	tnXstpPortStatsTable	v1.4
78	-THERMAL-PROTECTION-MIB	tnThermalProtectionPriorityTable	v1.4
79	-THERMAL-PROTECTION-MIB	tnThermalProtectionIfTable	v1.4
80	-THERMAL-PROTECTION-MIB	tnThermalProtectionIfStatusTable	v1.4
81	-ACCESS-MGMT-MIB	tnAccessMgmtCfgTable	v1.4
82	-ACCESS-MGMT-MIB	tnAccessMgmtTable	v1.4
83	-ACCESS-MGMT-MIB	tnAccessMgmtStatsTable	v1.4
84	-IP-SOURCE-GUARD-MIB	tnIPSourceGuardTable	v1.4
85	-IP-SOURCE-GUARD-MIB	tnIPSourceGuardIfTable	v1.4
86	-IP-SOURCE-GUARD-MIB	tnIPSourceGuardStaticTable	v1.4
87	-IP-SOURCE-GUARD-MIB	tnIPSourceGuardDynamicTable	v1.4
88	-ARP-INSPECTION-MIB	tnARPInspectionConfigTable	v1.4
89	-ARP-INSPECTION-MIB	tnARPInspectionPortModeTable	v1.4
90	-ARP-INSPECTION-MIB	tnStaticARPInspectionTable	v1.4
91	-ARP-INSPECTION-MIB	tnDynamicARPInspectionTable	v1.4
92	SNMP-USER-BASED-SM-MIB	usmUser	v1.4
93	SNMP-VIEW-BASED-ACM-MIB	vacmSecurityToGroupTable	v1.4

94	SNMP-VIEW-BASED-ACM-MIB	vacmAccessTable	v1.4
95	SNMP-VIEW-BASED-ACM-MIB	vacmMIBViews	v1.4
96	-DEV-VLAN-TRANSITION-MIB	tnVlanTransPort1GroupMapTable	v1.4
97	-DEV-VLAN-TRANSITION-MIB	tnVlanTransMapTable	v1.4
98	SNMP-TARGET-MIB	snmpTargetAddrTable	v1.4
99	SNMP-TARGET-MIB	snmpTargetParamsTable	v1.4
100	SNMP-NOTIFICATION-MIB	snmpNotifyTable	v1.4
101	SNMP-COMMUNITY-MIB	snmpCommunityTable	v1.4
102	SNMP-COMMUNITY-MIB	snmpTargetAddrExtTable	v1.4
103	DOT3-OAM-MIB	dot3OamEventLogTable	v1.4
104	-SECURITY-AAA-MIB	tnAAAServerTable	v1.4
105	-SECURITY-AAA-MIB	tnStatisticsTable	v1.4
106	-DEV-SYS-IPMGMT-MIB.smi	tnIPv4MgmtTable	v1.4
107	-DEV-SYS-IPMGMT-MIB.smi	tnIPv6MgmtTable	v1.4

Public MIBs

MIB	Table	Version
RFC1213-MIB	system	1.2.4
RFC1213-MIB	sysORTable	1.2.4
RFC1213-MIB	interfaces	1.2.4
RFC1213-MIB	ifTable	1.2.4
RFC1213-MIB	ip	1.2.4
RFC1213-MIB	ipAddrTable	1.2.4
RFC1213-MIB	ipNetToMediaTable	1.2.4
RFC1213-MIB	ip	1.2.4
RFC1213-MIB	icmp	1.2.4
RFC1213-MIB	tcp	1.2.4
RFC1213-MIB	tcpConnTable	1.2.4
RFC1213-MIB	udp	1.2.4
RFC1213-MIB	udpTable	1.2.4
RMON2-MIB	etherStatsTable	1.2.4
RMON2-MIB	historyControlTable	1.2.4
RMON2-MIB	alarmTable	1.2.4
RMON2-MIB	1.2.4	1.2.4
IEEE8021-PAE-MIB	1.2.4	1.2.4
IEEE8021-PAE-MIB	1.2.4	1.2.4
IEEE8021-PAE-MIB	1.2.4	1.2.4
Q-BRIDGE-MIB	dot1qBase	1.3.4
Q-BRIDGE-MIB	dot1qFdbTable	1.3.4
Q-BRIDGE-MIB	dot1qTpFdbTable	1.3.4
Q-BRIDGE-MIB	dot1qTpGroupTable	1.3.4
Q-BRIDGE-MIB	dot1qStaticUnicastTable	1.3.4
Q-BRIDGE-MIB	dot1qStaticMulticastTable	1.3.4
Q-BRIDGE-MIB	dot1qVlanCurrentTable	1.3.4
Q-BRIDGE-MIB	dot1qVlanStaticTable	1.3.4
Q-BRIDGE-MIB	dot1qPortVlanTable	1.3.4
IEEE8021-Q-BRIDGE-MIB	ieee8021QBridgeTable	1.3.4
IEEE8021-Q-BRIDGE-MIB	ieee8021QBridgeFdbTable	1.3.4
IEEE8021-Q-BRIDGE-MIB	ieee8021QBridgeTpFdbTable	1.3.4
IEEE8021-Q-BRIDGE-MIB	ieee8021QBridgeTpGroupTable	1.3.4
IEEE8021-Q-BRIDGE-MIB	ieee8021QBridgeStaticUnicastTable	1.3.4
IEEE8021-Q-BRIDGE-MIB	ieee8021QBridgeStaticMulticastTable	1.3.4
IEEE8021-Q-BRIDGE-MIB	ieee8021QBridgeVlanCurrentTable	1.3.4
IEEE8021-Q-BRIDGE-MIB	ieee8021QBridgeVlanStaticTable	1.3.4
IEEE8021-Q-BRIDGE-MIB	ieee8021QBridgeNextFreeLocalVlanTable	1.3.4
IEEE8021-Q-BRIDGE-MIB	ieee8021QBridgePortVlanTable	1.3.4
IEEE8021-SPANNING-TREE-MIB	ieee8021SpanningTreeTable	1.3.8
IEEE8021-SPANNING-TREE-MIB	ieee8021SpanningTreePortTable	1.3.8
IEEE8021-MSTP-MIB	ieee8021MstpCistTable	1.3.8
IEEE8021-MSTP-MIB	ieee8021MstpCistPortTable	1.3.8
IEEE8021-MSTP-MIB	ieee8021MstpTable	1.3.8
IEEE8021-MSTP-MIB	ieee8021MstpPortTable	1.3.8
IEEE8021-MSTP-MIB	ieee8021MstpConfigIdTable	1.3.8
IEEE8021-MSTP-MIB	ieee8021MstpVlanTable	1.3.8

LLDP-MIB	IldpPortConfigTable	1.3.8
LLDP-MIB	IldpConfigManAddrTable	1.3.8
LLDP-MIB	IldpStatsTxPortTable	1.3.8
LLDP-MIB	IldpStatsRxPortTable	1.3.8
LLDP-MIB	IldpLocalSystemData	1.3.8
LLDP-MIB	IldpLocPortTable	1.3.8
LLDP-MIB	IldpLocManAddrTable	1.3.8
LLDP-MIB	IldpRemTable	1.3.8
LLDP-MIB	IldpRemManAddrTable	1.3.8
IEEE8023-LAG-MIB	dot3adAggTable	1.3.8
IEEE8023-LAG-MIB	dot3adAggPortListTable	1.3.8
IEEE8023-LAG-MIB	dot3adAggPortTable	1.3.8
IEEE8023-LAG-MIB	dot3adAggPortStatsTable	1.3.8
IEEE8023-LAG-MIB	lagMIBObjects	1.3.8
SNMP-USER-BASED-SM-MIB	usmUser	1.3.10
SNMP-VIEW-BASED-ACM-MIB	vacmSecurityToGroupTable	1.3.10
SNMP-VIEW-BASED-ACM-MIB	vacmAccessTable	1.3.10
SNMP-VIEW-BASED-ACM-MIB	vacmMIBViews	1.3.10
SNMP-TARGET-MIB	snmpTargetAddrTable	
SNMP-TARGET-MIB	snmpTargetParamsTable	
SNMP-NOTIFICATION-MIB	snmpNotifyTable	
SNMP-COMMUNITY-MIB	snmpCommunityTable	
SNMP-COMMUNITY-MIB	snmpTargetAddrExtTable	

SNMP traps notes:

The Last Gasp can be in the form of IEEE802.3 2008 Clause 57 Dying gasp event and/or an SNMP trap to NMS system.

The Y.1731 AIS and LCK faults for fault monitoring and isolation raise SNMP traps.

All CCM errors such as remoteCCM, RDI, MACStatus, errorCCM, crossConnect, etc. are reported in MEP status and SNMP traps are raised for errors.

SNMP traps are generated for various Threshold events (Errored Symbol Period, Errored Frame Event, Errored Frame Period Event and Errored Frame Seconds summary events) and Non-threshold events (dying gasp and critical events).

The Link Fault, Dying Gasp, Critical Event, and other LOAM details for transmit and receive on each port are displayed at the **Monitor > Link OAM > Statistics** menu path.

Trap Primary or Secondary Power Supply: with both slot 1 and slot 2 connected, LED S2 is green and slot 1 is Primary. If slot 1 is removed, LED S2 becomes AMBER indicating that the LIB-44xx is operating with a Secondary power supply and will send out the trap as "entConfigChange".

For Additional MIB Information

For the list of LIB-44xx SNMP Traps see “[SNMP v3 Traps](#) on page 46.

For information on Link OAM MIB Retrieval see [Diagnostics > Link OAM](#) on page 412.

See the related sections of this manual for LIB-44xx configuration, monitoring, diagnostics, and maintenance via the LIB-44xx web interface (menu system). See the LIB-44xx CLI Reference manual for LIB-44xx configuration, monitoring, diagnostics, and maintenance via the CLI (Command Line Interface).

Glossary

This section describes many of the terms and mnemonics used in this manual. Note that the use of or description of a term does not in any way imply support of that feature or of any related function(s).

1+1

The Protection Type 1+1 uses the protection resources at all times for sending a replica of the traffic. The protection merge point, where both copies are expected to arrive, decides which of the two copies to select for forwarding.

The decision can be to switch from one resource to the other due to an event like resource up/down etc. or can be on a per frame/cell basis, the selection decision is performed according to parameters defined below (e.g. revertive, non-revertive, manual, etc.).

A network can offer protection by providing alternative resources to be used when the working resource fails.

The specific terminology for the number and arrangement of such resources includes 1+1, 1:1, 1:n, n:1, and m:n.

1:1

The 1:1 Protection Type provides a protection resource for a single working resource.

A network can offer protection by providing alternative resources to be used when the working resource fails.

The terminology for the number and arrangement of such resources includes 1+1, 1:1, 1:n, n:1, and m:n.

1 PPS

In IEEE 1588v2, a pulse that is repeated every second and has a very accurate phase. It synchronizes several geographically dispersed clients (e.g., cell sites) to the same time and phase of 1 μ s. Any third party test equipment must also support 1 PPS.

A

AAA

(Authentication, Authorization and Accounting); examples of this type of protocols include RADIUS, TACACS, TACACS+, etc. See the IETF Working Group [status](http://tools.ietf.org/wg/aaa/) page (<http://tools.ietf.org/wg/aaa/>) for more information. For IETF RFC information see <http://tools.ietf.org/html/rfc2975>.

Authentication: refers to the process where an entity's identity is authenticated, typically by providing evidence that it holds a specific digital identity such as an identifier and the corresponding credentials. Examples of types of credentials are passwords, one-time tokens, digital certificates, and phone numbers (calling/called).

Authorization: determines whether a particular entity is authorized to perform a given activity, typically inherited from authentication when logging on to an application or service. Authorization may be determined based on a range of restrictions, for example time-of-day restrictions, or physical location restrictions, or restrictions against multiple access by the same entity or user. Typical authorization in everyday computer life is for example granting read access to a specific file for authenticated user. Examples of types of service include IP address filtering, address assignment, route assignment, quality of Service/differential services, bandwidth control/traffic management, compulsory tunnelling to a specific endpoint, and encryption.

Accounting: refers to the tracking of network resource's consumption by users for the purpose of capacity and trend analysis, cost allocation, billing.[3] In addition, it may record events such as

authentication and authorization failures, and include auditing functionality, which permits verifying the correctness of procedures carried out based on accounting data. Real-time accounting refers to accounting information that is delivered concurrently with the consumption of the resources. Batch accounting refers to accounting information that is saved until it is delivered at a later time. Typical information gathered includes the identity of the user or other entity, the nature of the service delivered, when the service began, when it ended, and if there is a status to report.

ACE

ACE (**A**ccess **C**ontrol **E**ntry) describes the 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 application.

ACL

ACL is an acronym for **A**ccess **C**ontrol **L**ist. It is the list table of ACEs, containing access control entries that specify individual users or groups permitted or denied to 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 the various situation. 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 to use the service. ACL can generally be configured to control inbound traffic, and in this context, they are similar to firewalls.

There are three LIB-44xx web pages associated with the manual ACL configuration:

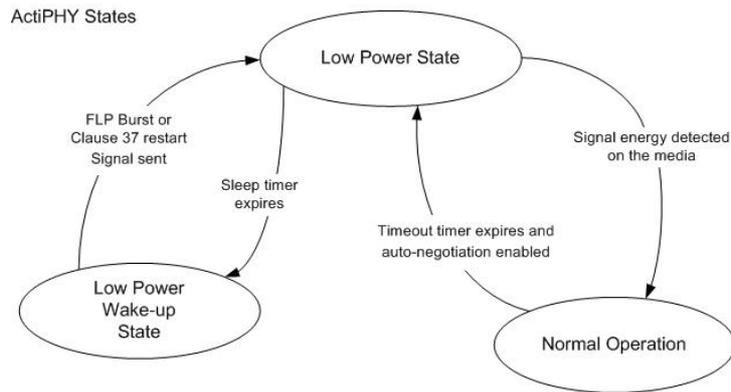
ACL|Access Control List: The web page shows the ACEs in a prioritized way, highest (top) to lowest (bottom). Default the table is empty. An ingress frame will only get a hit on one ACE even though there are more matching ACEs. The first matching ACE will take action (permit/deny) on that frame and a 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 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 seconds. Under "Ports" and "Access Control List" web-pages you can assign a Rate Limiter ID to the ACE(s) or ingress port(s).

ActiPHY™

An automatic power savings mode when a specific port is in link down or standby operation.



ActiPHY® is a registered trademark used for Semiconductors, Integrated Circuits and Ethernet Transceivers and owned by Vitesse Semiconductor Corporation.

Address

Digital information that uniquely identifies a network, station, device, etc. so that each can send and receive messages. There are four types of addresses commonly used with the Internet:

Email address (e.g., *name@mail_server.domain*)

IP address or Internet address: *a.b.c.d* or *device_name.sub-domain.domain*

MAC address (hardware address)

URL (Uniform Resource Locator): *method://server_adress[port]/document_path*

Address

In IPv6, an IPv6-layer identifier for an interface or a set of interfaces.

Alarm

The term 'alarm' actually refers to all types of fault events that are associated with a potential failure. Per MEF 15, the Perceived Alarm Severity (critical, major, minor, warning, indeterminate, or cleared).

Severity

assignments are only required for equipment alarms and physical layer communications alarms generated by the ME-NE).

- Critical - Indicates that a service affecting condition has occurred and immediate corrective action is required. Such a severity is used when the managed entity is totally out of service and its capability must be restored.
- Major - Indicates that a service affecting condition has occurred and urgent corrective action is required. Such a severity is used when there is a severe degradation in the capability of the managed entity and its full capability must be restored.
- Minor - Indicates that a non-service affecting condition has occurred and that corrective action should be taken in order to prevent a more serious fault.
- Warning - Indicates the detection of a potential or impending service affecting fault, before any significant effects have been felt.
- Indeterminate - The severity level cannot be determined.
- Cleared - The clearing of one or more previously reported alarms.

Anycast address

In IPv6, an identifier for a set of interfaces (typically belonging to different nodes). A packet sent to an anycast address is delivered to one of the interfaces identified by that address (the "nearest" one, according to the routing protocol's measure of distance).

AES

AES is an acronym for **A**dvanced **E**ncryption **S**tandard. The encryption key protocol is applied in 802.11 standard to improve WLAN security. It is an encryption standard by the U.S. government, which will replace DES and 3DES. AES has a fixed block size of 128 bits and a key size of 128, 192, or 256 bits.

AP

Access Point, such as a wireless Access Point defined by IEEE 802.11.

APS

APS is an acronym for **A**utomatic **P**rotection **S**witching. This protocol is used to secure that switching is done bidirectional in the two ends of a protection group, as defined in G.8031.

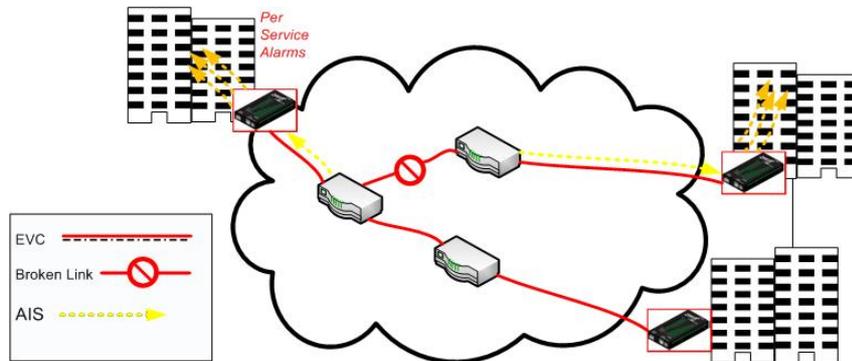
Aggregation

Using multiple ports in parallel to increase the link speed beyond the limits of a port and to increase the redundancy for higher availability. (Also *Port [Aggregation](#), Link Aggregation*).

Alarm Indication Signal (AIS)

ETH-AIS allows alarm suppression when defects are to be detected at the server layer. You can enable or disable frames transmission with ETH-AIS information on an MEP or on a server MEP. You can also issue frames with ETH-AIS information at the client maintenance level by a MEP, including a server MEP, on detecting defect conditions. Defect conditions can include signal fail conditions with ETH-CC enabled, and AIS condition with ETH-CC disabled.

Only a MEP or Server MEP is configured to issue frames with ETH-AIS information. When a MEP detects a defect condition, it immediately starts transmitting periodic frames with ETH-AIS information at a configured client maintenance level. The MEP continues to transmit periodic frames with ETH-AIS information until the defect condition is resolved. On receiving a frame with ETH-AIS information, a MEP detects the AIS condition and suppresses loss of continuity alarms with all of its peer MEPs. The MEP resumes loss of continuity alarm generation on detecting loss of continuity conditions in place of the AIS condition.



Transmission of frames with ETH-AIS information can be enabled or disabled on a MEP (or on a server MEP). Frames with ETH-AIS information can be issued at the client MEG level by a MEP, including a server MEP, upon detecting defect conditions.

ARP

ARP is an acronym for **A**ddress **R**esolution **P**rotocol. It is a protocol that used to convert an [IP](#) address into a physical address, such as an Ethernet address. ARP allows a host to communicate with other hosts when only the Internet address of its neighbours is known. Before using IP, the host sends a broadcast ARP request containing the Internet address of the desired destination system.

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.

Automatic Reversion

The protection is in revertive mode if, after a resource failure and its subsequent repair, the network automatically reverts to using this initial resource. The protection is in non-revertive mode otherwise. Automatic reversion may include a reversion timer (i.e., the Wait To Restore), which delays the time of reversion after the repair.

Auto-Negotiation

Auto-negotiation is the process where two different devices establish the mode of operation and the speed settings that can be shared by those devices for a link.

B

Bandwidth Profile

A characterization of ingress Service Frame arrival times and lengths at a reference point and a specification of the disposition of each Service Frame based on its level of compliance with the Bandwidth Profile. In MEF documents, the reference point is the UNI. See [MEF 6.1](#).

Boundary clock

A clock that has multiple Precision Time Protocol (PTP) ports in a domain and maintains the timescale used in the domain. It may serve as the source of time (i.e., be a master clock) and may synchronize to another clock (i.e., be a slave clock).

A boundary clock is a clock with more than a single PTP port, with each PTP port providing access to a separate PTP communication path. Boundary clocks are used to eliminate fluctuations produced by routers and similar network elements.

Broadcast

A message forwarded to all (multiple, unspecified recipients) network destinations. On Ethernet, a broadcast packet is a special type of multicast packet where all nodes on the network are always willing to receive.

C

CAPWAP

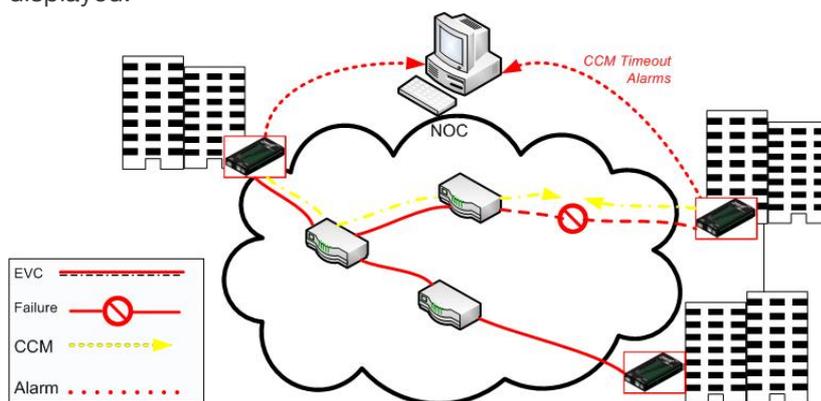
Control and Provisioning of Wireless Access Points.

CC

CC (Continuity Check) is a [MEP](#) function that detects loss of continuity in a network by transmitting [CCM](#) frames to a peer MEP.

CC Monitoring (Continuity Checks Monitoring)

Fault detection uses the Continuity Check protocol to detect both connectivity failures and unintended connectivity between service instances. Each MEP can periodically transmit a multicast Connectivity Check Message (CCM) announcing the identity of the MEP and its MA, and tracks the CCMs received from the other MEPs. All connectivity faults that can misdirect a CCM show up as differences between the CCMs received and the MEP's configured expectations. The state of the tracked CCMs can be displayed.



Each Continuity Check Message (CCM) is a multicast CFM PDU transmitted periodically by a MEP to ensure continuity over the MA to which the transmitting MEP belongs. No reply is sent by any MP in response to receiving a CCM. CCMs use addresses from the Continuity Check Message Group Destination MAC Address table. The CCM can be sent away from or towards the MAC Relay Entity.

CCM

CCM is an acronym for Continuity Check Message. It is a [OAM](#) frame transmitted from a MEP to its peer MEP and used to implement [CC](#) functionality.

CDP

CDP (Cisco Discovery Protocol) is a Cisco proprietary Layer 2 protocol that is media- and protocol-independent, and runs on Cisco routers, bridges, access servers, and switches. A Cisco device with CDP enabled sends out periodic interface updates to a multicast address in order to make itself known to neighbours. As a layer two protocol, these packets (frames) are not routed. Using SNMP with the CDP MIB lets network management applications learn the device type and the SNMP agent address of neighbouring devices, and to then send SNMP queries to those devices.

CIST

Acronym for **C**ommon and **I**nternal **S**panning **T**ree. Concerning IST/CST/CIST, IST is the only instance that can send and receive BPDUs in the MST network. An MSTn instance is local to a region. ISTs in different regions are interconnected via a Common Spanning Tree (CST). The CIST includes the collection of ISTs in each MST region, and the CST that connects the ISTs.

The CIST is the default spanning tree instance of MSTP (i.e., all VLANs that are not members of particular MSTIs are members of the CIST. Also, an individual MST region can be regarded a single virtual bridge by other MST regions. The spanning tree that runs between MSTP regions is the CIST.

Clock

In PTP, a node participating in the Precision Time Protocol (PTP) that is capable of providing a measurement of the passage of time since a defined epoch.

Commonly Used EtherTypes

The 'EtherType' field in an Ethernet frame indicates the protocol used in the data field of the frame. According to the IEEE 802.3, Length/EtherType field is a two-octet field which takes one of two meanings, depending on its numeric value. For numeric evaluation, the first octet is the most significant octet; when the value of this field is ≥ 1536 decimal (0600 hex) the EtherType field indicates the nature of the MAC client protocol (EtherType interpretation). The value of the Type Field is obtained from the IEEE EtherType Field Registrar. The EtherType field is a very limited space and assignments are limited. The EtherType field is administered by the IEEE RAC EtherType Field Approval Authority. The following list of EtherTypes is unverified information from various sources.

EtherType (hex)	Protocol
0x000 - 0x05DC	IEEE 802.3 length
0x0101- 0x01FF	Experimental
0x0600	Xerox NS IDP
0x0660, 0x0661	DLOG
0x0800	IP (Internet Protocol)
0x0801	X.75 Internet
0x0802	NBS Internet
0x0803	ECMA Internet
0x0804	Chaosnet
0x0805	X.25 Level 3
0x0806	ARP (Address Resolution Protocol)
0x0808	Frame Relay ARP (RFC 1701)
0x6559	Raw Frame Relay (RFC 1701)
0x8035	RARP (Reverse Address Resolution Protocol), DRAP (Dynamic RARP)
0x8037	Novell Netware IPX
0x809B	EtherTalk (AppleTalk over Ethernet)
0x80D5	IBM SNA Services over Ethernet
0x80F3	AARP, AppleTalk Address Resolution Protocol

0x8100	VLAN-tagged frame (IEEE 802.1Q)
0x8137	IPX (Internet Packet Exchange)
0x814c	SNMP (Simple Network Management Protocol)
0x86DD	IPv6 (Internet Protocol version 6)
0x8808	MAC Control
0x8809	Slow Protocols (IEEE 802.3)
0x880B	PPP (Point to Point Protocol)
0x880C	GSMP (General Switch Management Protocol)
0x8819	CobraNet
0x8847	MPLS (Multi-Protocol Label Switching) (unicast)
0x8848	MPLS (Multi-Protocol Label Switching) (multicast)
0x8863	PPoE (PPP over Ethernet) (Discovery stage)
0x8864	PPoE (PPP over Ethernet) (PPP Session stage)
0x886F	Microsoft NLB heartbeat
0x8870	Jumbo Frames
0x887B	HomePlug 1.0 MME
0x888E	EAPOL (EAP over LAN) (IEEE 802.1X)
0x88BB	LWAP (Light Weight Access Point Protocol)
0x88CC	LLDP (Link Layer Discovery Protocol)
0x8892	PROFINET Protocol
0x889A	HyperSCSI (SCSI over Ethernet)
0x88A2	ATA over Ethernet
0x88A4	EtherCAT Protocol
0x88A8	Provider Bridging (IEEE 802.1ad)
0x88AB	Ethernet Powerlink
0x88CC	LLDP
0x88CD	SERCOS III
0x88D8	Circuit Emulation Services over Ethernet (MEF-8)
0x88E1	HomePlug AV MME
0x88E5	MAC security (IEEE 802.1AE)
0x88F7	Precision Time Protocol (IEEE 1588)
0x8902	IEEE 802.1ag Connectivity Fault Management (CFM) Protocol / ITU-T Recommendation Y.1731 (OAM)
0x8906	Fibre Channel over Ethernet
0x8914	FCoE Initialization Protocol
0x9000	Loopback (Configuration test protocol)
0x9100	VLAN Tag Protocol Identifier (Q-in-Q)
0x9200	VLAN Tag Protocol Identifier
0xCAFE	Veritas LLT (Low Latency Transport)
0xFFFF	(Reserved)

Note: Some well known EtherTypes are not necessarily listed in the IEEE list of EtherType values. For example, EtherType 0x0806 (used by ARP) is listed by the IEEE only as "Symbolics, Protocol unavailable."

See <http://standards.ieee.org/develop/regauth/ethertype/eth.txt> for more information.

The EtherType is one of two types of protocol identifier parameters that can occur in Ethernet frames after the initial MAC-48 destination and source identifiers. Ethertypes are 16-bit identifiers appearing as the initial two octets after the MAC destination and source (or after a tag).

EtherType use implies the use of the IEEE Assigned EtherType Field with IEEE Std 802.3, 1998 Edition Local and Metropolitan Area Networks. The EtherType Field provides a context for interpretation of the

data field of the frame (protocol identification). Several well-known protocols already have an EtherType Field.

The IEEE 802.3, 1998 Length/EtherType Field, originally known as EtherType, is a two-octet field. When the value of this field is greater than or equal to 1536 decimal (0600 hexadecimal) the EtherType Field indicates the nature of the MAC client protocol (EtherType interpretation). The length and EtherType interpretations of this field are mutually exclusive.

Communication

In IPv6, any packet exchange among nodes that requires that the address of each node used in the exchange remain the same for the duration of the packet exchange. Examples are a TCP connection or a UDP request- response.

CoS

The QoS technique known as Class of Service (CoS) is a 3-bit field called the Priority Code Point (PCP) within an Ethernet frame header when using VLAN tagged frames as defined by IEEE 802.1Q. The PCP specifies a priority value of between 0 and 7 (inclusive) to be used by QoS disciplines to differentiate traffic. This technique is commonly referred to as IEEE 802.1p, but there is no IEEE standard or amendment under that name; the technique is incorporated into the IEEE 802.1Q standard, which specifies the tag inserted into an Ethernet frame.

Eight different classes of service can be expressed with the 3-bit PCP field in an IEEE 802.1Q header added to the frame. The way traffic is treated when assigned to any particular class is undefined by the spec and is left to the implementation. The IEEE however has made some broad recommendations:

<u>PCP</u>	<u>Priority</u>	<u>Acronym</u>	<u>Traffic Types</u>
1	0 (lowest)	BK	Background
0	1	BE	Best Effort
2	2	EE	Excellent Effort
3	3	CA	Critical Applications
4	4	VI	Video, < 100 ms latency and jitter
5	5	VO	Voice, < 10 ms latency and jitter
6	6	IC	Internetwork Control
7	7 (highest)	NC	Network Control

Note: the above recommendation was revised in IEEE 802.1Q-2005, and it also differs from the original IEEE 802.1D-2004 recommendation. See also "QoS".

D

DA

(Destination Address); contrast SA.

DAD

(Duplicate Address Detection) - In IPv6, part of the NDP protocol that lets nodes check if an address is already in use.

DEI

[DEI](#) is an acronym for **D**rop **E**ligible **I**ndicator. It is a 1-bit field in the VLAN tag.

Deprecated address

In IPv6, an address assigned to an interface whose use is discouraged, but not forbidden. A deprecated address should no longer be used as a source address in new communications, but packets sent from or to deprecated addresses are delivered as expected. A deprecated address may continue to be used as a source address in communications where switching to a preferred address causes hardship to a specific upper-layer activity (e.g., an existing TCP connection).

DES

DES (**D**ata **E**ncryption **S**tandard) provides a complete description of a mathematical algorithm for encrypting (enciphering) and decrypting (deciphering) binary coded information.

Encrypting data converts it to an unintelligible form called cipher. Decrypting cipher converts the data back to its original form called plaintext. The algorithm described in this standard specifies both enciphering and deciphering operations which are based on a binary number called a key.

DHCP

DHCP (Dynamic Host Configuration Protocol) is a protocol used for assigning dynamic IP addresses to devices on a network. DHCP used by networked computers (clients) to obtain IP addresses and other parameters such as the default gateway, subnet mask, and IP addresses of DNS servers from a DHCP server. The DHCP server ensures that all IP addresses are unique, for example, no IP address is assigned to a second client while the first client's assignment is valid (its lease has not expired). Therefore, IP address pool management is done by the server and not by a human network administrator.

Dynamic addressing simplifies network administration because the software keeps track of IP addresses rather than requiring an administrator to manage the task. This means that a new computer can be added to a network without the hassle of manually assigning it a unique IP address.

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.

The DHCP option 82 enables a DHCP relay agent to insert specific information into a DHCP request packets when forwarding client DHCP packets to a DHCP server and remove the specific information from a DHCP reply packets when forwarding server DHCP packets to a DHCP client. The DHCP server can use this information to implement IP address or other assignment policies. Specifically the option works by setting two sub-options: Circuit ID (option 1) and Remote ID (option2). The Circuit ID sub-option is supposed to include information specific to which circuit the request came in on. The Remote ID sub-option was designed to carry information relating to the remote host end of the circuit.

The definition of Circuit ID in the switch is 4 bytes in length and the format is "vlan_id" "module_id" "port_no". The parameter of "vlan_id" is the first two bytes represent the VLAN ID. The parameter of "module_id" is the third byte for the module ID (in standalone switch it always equal 0). The parameter of "port_no" is the fourth byte and it means the port number. The Remote ID is 6 bytes in length, and the value is equal the DHCP relay agents MAC address.

DHCP Snooping

DHCP Snooping is used to block intruders 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 DHCP client and server.

DMAC

(Destination MAC Address) A valid source MAC address, except for an address which has the lowest bit of the first byte set to '1'. These addresses, including the all 1's broadcast address FF:FF:FF:FF:FF:FF and the set of multicast addresses, are point-to-multipoint addresses and can never appear as the source address in an Ethernet frame. Note that a frame must be sent by a single source.

Each MAC header consists of three parts: 1. A 6-byte destination address, which specifies either a single recipient node (unicast mode), a group of recipient nodes (multicast mode), or the set of all recipient nodes (broadcast mode). 2. A 6-byte source address, which is set to the sender's globally unique node address. This may be used by the network layer protocol to identify the sender, but usually other mechanisms are used (e.g. ARP). Its main function is to allow address learning which may be used to configure the filter tables in a bridge.

3. A 2-byte type field, which provides a Service Access Point (SAP) to identify the type of protocol being carried.

See also "SMAC".

DMI

Diagnostic Monitoring Interface; the LIB-44xx is capable of supporting connectors with DMI (SFF-8472) capability. All DMI events will trigger notification. An intrusion detection based on Rx Power level is available for triggering any drop in the Rx power.

DNS

DNS (Domain Name System) stores and associates many types of information with domain names. Most importantly, DNS translates human-friendly domain names and computer hostnames into computer-friendly IP addresses. For example, the domain name www.example.com might translate to 192.168.0.1.

DoS

DoS is an acronym for Denial of Service. In a DoS attack, an attacker attempts to prevent legitimate users from accessing information or services. By targeting at network sites or network connection, an attacker may be able to prevent network users from accessing email, web sites, online accounts (banking, etc.), or other services that rely on the affected computer.

Dotted Decimal Notation

Dotted Decimal Notation refers to a method of writing IP addresses using decimal numbers and dots as separators between octets. An IPv4 dotted decimal address has the form x.y.z.w, where x, y, z, and w are decimal numbers between 0 and 255.

DSCP

DSCP (Differentiated Services Code Point) is a field in the header of [IP](#) packets for packet classification purposes. In an IP header, a six-bit DSCP field specifies the per-hop behaviour for a given flow of packets. Each packet is given one of 64 possible forwarding behaviours (known as per-hop behaviours, or PHBs) for a given set of packet travel rules. DSCP uses the first 6 bits in the ToS field of the IPv4 packet header. In many cases, DSCP has replaced the outdated Type of Service (TOS) field.

Dual stack

One of three options for migrating to IPv6 from an existing IPv4 network infrastructure (dual-stack network, tunnelling, and translation).

E**E911**

Enhanced 911 Emergency Call Service applicable in North America.

EAPOL

The key protocol in 802.1x is called 'EAP over LANs' (EAPOL), which is currently defined for Ethernet-like LANs including 802.11 wireless, as well as token ring LANs (including FDDI).

In 802.1X, the user is called the 'supplicant', the switch is the 'authenticator', and the RADIUS server is the 'authentication server'. The authenticator acts as the man-in-the-middle, forwarding requests and responses between the supplicant and the authentication server. Frames sent between the supplicant and the switch are special 802.1X frames, known as EAPOL (EAP Over LANs) frames. EAPOL frames encapsulate EAP PDUs (RFC3748). Frames sent between the switch and the RADIUS server are RADIUS packets. RADIUS packets also encapsulate EAP PDUs together with other attributes like the switch's IP address, name, and the supplicant's port number on the switch. EAP is very flexible, in that it allows for different authentication methods, like MD5-Challenge, PEAP, and TLS. Note that the authenticator (the switch) doesn't need to know which authentication method the supplicant and the authentication server are using, or how many information exchange frames are needed for a particular method. The switch simply encapsulates the EAP part of the frame into the relevant type (EAPOL or RADIUS) and forwards it.

The authenticator acts like a 'security guard' to a protected network. The supplicant (client device) is not allowed access through the authenticator to the protected side of the network until the supplicant's identity has been validated and authorized. The commonly used EtherType for EAPOL is 0x888E.

ECS

Emergency Call Service (e.g. E911 and others), such as defined by TIA or NENA.

ECEs

(EVC Control Entries) The ECE ID identifies the ECE. Unique ECE IDs are automatically assigned to ECEs added. The possible range is from 1 through 128. See also "EVC".

EEE

EEE is an abbreviation for Energy Efficient Ethernet defined in IEEE 802.3az.

ELIN

Emergency Location Identification Number, a valid North America Numbering Plan format telephone number, supplied to the PSAP for ECS purposes.

ENNI

(External Network-to-Network Interface) External Network to Network Interface; a reference point representing the boundary between two Operator MENs that are operated as separate administrative domains per MEF 26, 30. Previously "E-NNI".

Epoch

The origin of a PTP timescale.

EPS / ELPS

EPS is an abbreviation for Ethernet Protection Switching defined in ITU-T G.8031 (Ethernet (Linear) Protection Switch). Rec. ITU-T G.8031/Y.1342 (11/2009) defines the automatic protection switching APS protocol and linear protection switching mechanisms for point-to-point VLAN-based ETH SNC (Subnetwork Connection) in Ethernet transport networks. Protection switching occurs based on detection of certain defects on the transport entities (working and protection) within the protected domain. These defects are discussed in ITU-T G.8021.

The G.8031 Recommendation specifies linear protection switching mechanisms to be applied to VLAN-based Ethernet networks as described in G.8010. Protection switching is a fully allocated survivability mechanism ('fully allocated' in that the route and bandwidth of the protection entity is reserved for a selected working entity). EPS provides a fast and simple survivability mechanism. It is easier for a network operator to understand the network status (e.g., active network topology) with EPS than with other survivability mechanisms such as RSTP.

G.8031 specifies linear 1+1 protection switching architecture and linear 1:1 protection switching architecture.

The ETH-APS defined in Y.1731 is used as a signalling channel. [G.8031 \(2006\) Amd. 1 renamed EPS to ELPS.](#)

ERP instance

An entity that is responsible for the protection of a subset of the VLANs that transport traffic over the physical Ethernet ring. Each ERP instance is independent of other ERP instances that may be configured on the physical Ethernet ring. Per ITU-T Rec.G.8032/Y.1344 (03/2010).

ERPS

ERPS is an abbreviation for Ethernet ring protection switching. Recommendation ITU-T G.8032/Y.1344 defines the automatic protection switching (APS) protocol and protection switching mechanisms for ETH layer Ethernet ring topologies. Included are details on Ethernet ring protection characteristics and architectures, and the Ring APS (R-APS) protocol. The protection protocol defined in this Recommendation enables protected point-to-point, point-to-multipoint and multipoint-to-multipoint connectivity within a ring or interconnected rings, called "multi-ring/ladder network" topology. The ETH layer ring maps to the physical layer ring structure.

The ERPS effort at ITU-T under G.8032 is to provide sub-50ms protection and recovery switching for Ethernet traffic in a ring topology and at the same time ensuring that there are no loops formed at the Ethernet layer. G.8032v1 supported a single ring topology and G.8032v2 supports multiple rings/ladder topology.

ERPS specifies protection switching mechanisms and a protocol for Ethernet layer network (ETH) rings. Ethernet Rings can provide wide-area multipoint connectivity more economically due to their reduced number of links. The mechanisms and protocol defined in G.8032 provide highly reliable and stable protection; and avoid loops which would prove fatal to network operation and service availability.

Each Ethernet Ring Node is connected to adjacent Ethernet Ring Nodes participating in the same Ethernet Ring, using two independent links. A ring link is bounded by two adjacent Ethernet Ring Nodes, and a port for a ring link is called a ring port. The minimum number of Ethernet Ring Nodes in an Ethernet Ring is two.

The basis of this RPS architecture are a) the principle of loop avoidance, and b) the use of learning, forwarding, and Filtering Database (FDB) mechanisms defined in the Ethernet flow forwarding function (ETH_FF).

Loop avoidance in an Ethernet Ring is done by guaranteeing that at all times, 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 for blocking 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 Neighbour Node, may also participate in blocking or unblocking its end of the RPL.

An Ethernet Ring failure results in protection switching of the traffic. This is achieved under the control of the ETH_FF functions on all Ethernet Ring Nodes. An APS protocol is used to coordinate the protection actions over the ring.

ERPS Performance

Note from Rec. ITU-T G.8032/Y.1344 (03/2010): "Ethernet ring protection switching performance: 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 fibre 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 will be less than 50 ms. On Ethernet rings under all other conditions, the switch completion time may exceed 50 ms (the specific interval is under study), to allow time to negotiate and accommodate coexisting APS requests. In case of interconnection of sub-rings with R-APS virtual channel to a major ring, the R-APS messages of the

sub-ring that are inserted into the R-APS virtual channel take on performance characteristics (e.g., delay, jitter, packet drop probability, etc.) of the ring links and Ethernet ring nodes it crosses over the interconnected Ethernet ring. In this case, if the R-APS channel and R-APS virtual channel exceed the number of Ethernet ring nodes or fibre circumference defined above, the protection switching of the sub-ring may exceed 50 milliseconds. NOTE – The inclusion of the completion of FDB flush operation within the transfer time is for further study.”

ESP

The IP Encapsulating Security Payload (ESP) protocol provides a mix of security services in IPv4 and IPv6.

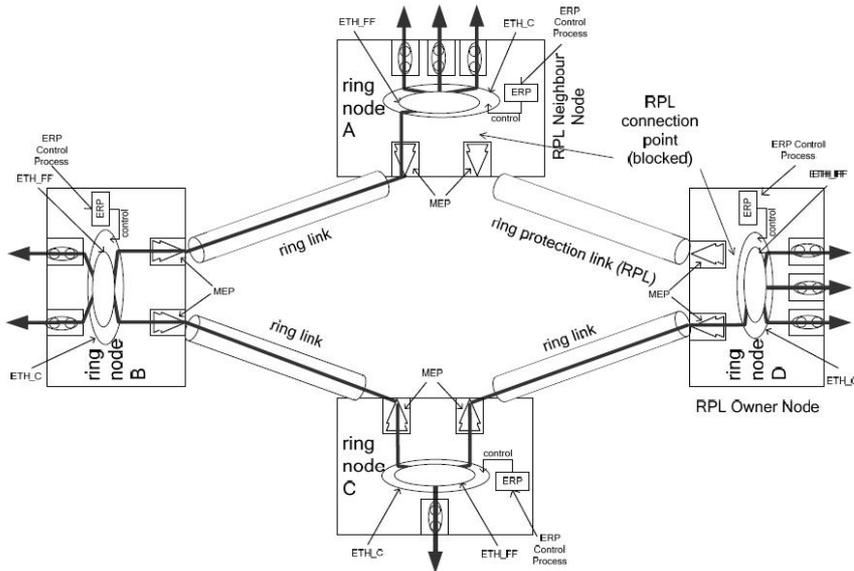
ESP supports two modes of operation: tunnel mode and transport mode.

The ESP header is designed to provide a mix of security services in IPv4 and IPv6. ESP may be applied alone, in combination with AH, or in a nested fashion.

Security services can be provided between a pair of communicating hosts, between a pair of communicating security gateways, or between a security gateway and a host. The ESP header is inserted after the IP header and before the next layer protocol header (transport mode) or before an encapsulated IP header (tunnel mode). ESP can be used to provide confidentiality, data origin authentication, connectionless integrity, an anti-replay service (a form of partial sequence integrity), and (limited) traffic flow confidentiality. The set of services provided depends on options selected at the time of Security Association (SA) establishment and on the location of the implementation in a network topology. See IETF [RFC 4303](#).

Ethernet ring

A collection of Ethernet ring nodes forming a closed physical loop whereby each Ethernet ring node is connected to two adjacent Ethernet ring nodes via a duplex communications facility. From ITU-T Rec.G.8032/Y.1344 (03/2010).



Ethernet ring node

A network element which implements at least the following functionalities:

- One Ethernet connection function (ETH_C) with a dedicated Ethernet flow forwarding function (ETH_FF) for forwarding ring automatic protection switching (R-APS) control traffic.
- Two ring ports, including ETHDi/ETH adaptation function at the ring maintenance entity group level (MEL).
- Ethernet ring protection (ERP) control process controlling the blocking and unblocking of traffic over the ring ports. Per ITU-T Rec.G.8032/Y.1344 (03/2010).

Ethernet Services

Generally refers to Metro Ethernet Services available from service providers (SPs) per MEF specifications (MEF 6, Ethernet Services Definitions, and MEF 10, Ethernet Services Attributes).

Ethernet Type

Ethernet Type, or EtherType, is a field in the Ethernet MAC header, defined by the Ethernet networking standard. It is used to indicate which protocol is being transported in an Ethernet frame. See “Commonly Used EtherTypes” above.

EUI-64

The 64-bit Extended Unique Identifier (EUI-64) in IPv6.

EVC

(Ethernet Virtual Connection) An association of two or more UNIs that limits the exchange of frames to UNIs in the EVC. Generally, an EVC allows Ethernet service frames to be exchanged between UNIs that are connected via the same EVC. Per MEF 6.1, EVC performance requires “At least one CoS is REQUIRED. MUST specify CoS ID, per section 6.8 of [2]. MUST list values for each of the following attributes {Frame Delay, Frame Delay Variation, Frame Loss Ratio, and Availability} for each CoS, where Not Specified (N/S) is an acceptable value.”

F**Fast Leave**

Multicast snooping [Fast Leave](#) processing allows the switch to remove an interface from the forwarding-table entry without first sending out group specific queries to the interface. The VLAN interface is “pruned” from the multicast tree for the multicast group specified in the original leave message. Fast leave is enabled at the port level. Pruning happens per VLAN per port. See the IGMP group table which is indexed by VLAN and group. Fast-leave processing ensures optimal bandwidth management for all hosts on a switched network, even when multiple multicast groups are in use simultaneously. This processing applies to IGMP and MLD.

FCS

(Frame Check Sequence) per MEF 8, 11, 12.

flow

A given type of traffic sent between a producer device through a network to a endpoint known as a consumer. As the traffic goes through the network it, “flows” through the network. See also "Per flow QoS".

Foreign master

An ordinary or boundary clock sending Announce messages to another clock that is not the current master recognized by the other clock.

FPGA

(Field-Programmable Gate Array) a chip that can be programmed in the field after manufacture.

FTP

FTP is an acronym for **F**ile **T**ransfer **P**rotocol. It is a transfer protocol that uses the Transmission Control Protocol ([TCP](#)) and provides file writing and reading. It also provides directory service and security features.

G**Global address**

In IPv6, an address with unlimited scope.

Grandmaster clock

Within a PTP domain, a clock that is the ultimate source of time for clock synchronization using the protocol.

H

HMAC

(Hash-based Message Authentication Code) - a specific construction that calculates a message authentication code (MAC) involving a cryptographic hash function in combination with a secret key. As with any MAC, it may be used to simultaneously verify both the data integrity and the authenticity of a message. Any cryptographic hash function (e.g., MD5 or SHA-1) may be used in the calculation of an HMAC; the resulting MAC algorithm is termed HMAC-MD5 or HMAC-SHA1 accordingly. The cryptographic strength of the HMAC depends on the cryptographic strength of the underlying hash function, the size of its hash output length in bits, and on the size and quality of the cryptographic key.

Host

In IPv6, any node that is not a router.

HPIC

(Helper PIC) a PIC specific to the LIB-44xx ET (external timing) board.

HTTP

HTTP (Hypertext Transfer Protocol) is a protocol that used to transfer or convey information on the World Wide Web (WWW). HTTP defines how messages are formatted and transmitted, and what actions Web servers and browsers should take in response to various commands. For example, when you enter a URL in your browser, this actually sends an HTTP command to the Web server directing it to fetch and transmit the requested Web page. The other main standard that controls how the World Wide Web works is HTML, which covers how Web pages are formatted and displayed.

Any Web server machine contains, in addition to the Web page files it can serve, an HTTP daemon, a program that is designed to wait for HTTP requests and handle them when they arrive. The Web browser is an HTTP client, sending requests to server machines. An HTTP client initiates a request by establishing a Transmission Control Protocol ([TCP](#)) connection to a particular port on a remote host (port 80 by default). An HTTP server listening on that port waits for the client to send a request message.

HTTPS

HTTPS (Hypertext Transfer Protocol over Secure Socket Layer) is used to indicate a secure [HTTP](#) connection. HTTPS provide authentication and encrypted communication and is widely used on the World Wide Web for security-sensitive communication such as payment transactions and corporate logons.

HTTPS is really just the use of Netscape's Secure Socket Layer (SSL) as a sublayer under its regular HTTP application layering. (HTTPS uses port 443 instead of HTTP port 80 in its interactions with the lower layer, [TCP/IP](#).) SSL uses a 40-bit key size for the RC4 stream encryption algorithm, which is considered an adequate degree of encryption for commercial exchange.

I

ICMP

ICMP (Internet Control Message Protocol) is a protocol that generated the error response, diagnostic or routing purposes. ICMP messages generally contain information about routing difficulties or simple exchanges such as time-stamp or echo transactions. For example, the [PING](#) command uses ICMP to test an Internet connection.

ICMPv6

(Internet Control Message Protocol version 6) is the implementation of the Internet Control Message Protocol (ICMP) for Internet Protocol version 6 (IPv6) defined in RFC 4443.[1] ICMPv6 is an integral part of IPv6 and performs error reporting, diagnostic functions (e.g., ping), and a framework for extensions to implement future changes. Several extensions are published to define new ICMPv6 message types and options for existing ICMPv6 message types. The Neighbour Discovery Protocol (NDP) is a node

discovery protocol in IPv6 that replaces and enhances functions of ARP. Secure Neighbour Discovery Protocol (SEND) is an extension of NDP with extra security. Multicast Router Discovery (MRD) allows discovery of multicast routers.

IEEE 802.1X

IEEE 802.1X is an IEEE standard for port-based Network Access Control. It provides authentication to devices attached to a LAN port, establishing a point-to-point connection or preventing access from that port if authentication fails. With 802.1X, access to all switch ports can be centrally controlled from a server, which means that authorized users can use the same credentials for authentication from any point within the network.

IGMP

IGMP (Internet Group Management Protocol) is a communications protocol used to manage the membership of Internet Protocol multicast groups. IGMP is used by IP hosts and adjacent multicast routers to establish multicast group memberships. It is an integral part of the IP multicast specification, like ICMP for unicast connections. IGMP can be used for online video and gaming, and allows more efficient use of resources when supporting these uses.

Enabling IGMP allows the ports to detect IGMP queries and report packets and manage IP multicast traffic through the switch. IGMP have three fundamental types of message as follows:

Message	Description
Query	A message sent from the querier (IGMP router or switch) asking for a response from each host belonging to the multicast group.
Report	A message sent by a host to the querier to indicate that the host wants to be or is a member of a given group indicated in the report message.
Leave Group	A message sent by a host to the querier to indicate that the host has quit to be a member of a specific multicast group.

IGMP Querier

When a router sends IGMP Query messages onto a particular link, this router is called the 'Querier'. In order for IGMP, and thus IGMP snooping, to function, a multicast router must exist on the network and generate IGMP queries. The tables created for snooping (holding the member ports for each multicast group) are associated with the querier. Without a querier the tables are not created and snooping will not work. Furthermore IGMP general queries must be unconditionally forwarded by all switches involved in IGMP snooping.[1] Some IGMP snooping implementations include full querier capability. Others are able to proxy and retransmit queries from the multicast router.

IGMP snooping

The process of listening to Internet Group Management Protocol (IGMP) network traffic. IGMP snooping, as implied by the name, is a feature that allows a network switch to listen in on the IGMP conversation between hosts and routers. By listening to these conversations, the switch maintains a map of which links need which IP multicast streams. Multicasts may be filtered from the links which do not need them. A switch will, by default, flood multicast traffic to all the ports in a broadcast domain (or the VLAN equivalent). Multicast can cause unnecessary load on host devices by requiring them to process packets they have not solicited. When purposefully exploited this is known as one variation of a denial-of-service attack. IGMP snooping is designed to prevent hosts on a local network from receiving traffic for a multicast group they have not explicitly joined. It provides switches with a mechanism to prune multicast traffic from links that do not contain a multicast listener (an IGMP client).

IGMP snooping allows a switch to only forward multicast traffic to the links that have solicited them. Essentially, IGMP snooping is a layer 2 optimization for the layer 3 IGMP. IGMP snooping takes place internally on switches and is not a protocol feature. Two standards organizations define IGMP snooping - the IEEE standardizes Ethernet switches, and the IETF standardizes IP multicast.

IMAP

IMAP (Internet Message Access Protocol) is a protocol for email clients to retrieve email messages from a mail server. IMAP is the protocol that IMAP clients use to communicate with the servers, and [SMTP](#) is the protocol used to transport mail to an IMAP server.

The current version of the Internet Message Access Protocol is IMAP4. It is similar to Post Office Protocol version 3 ([POP3](#)), but offers additional and more complex features. For example, the IMAP4 protocol leaves your email messages on the server rather than downloading them to your computer. If you wish to remove your messages from the server, you must use your mail client to generate local folders, copy messages to your local hard drive, and then delete and expunge the messages from the server.

I-NNI

(Internal Network to Network Interface) per MEF 4. Internal NNI (this definition has not been implemented in any specification) per MEF 17.

Interconnection node

An Ethernet ring node which is common to two or more Ethernet rings or to a sub-ring and an interconnected network. At each interconnection node there may be one or more Ethernet rings that can be accessed through a single ring port and not more than one Ethernet ring that is accessed by two ring ports. The former set of Ethernet rings is comprised of sub-rings, whereas the latter Ethernet ring is considered a major ring, relative to this interconnection node. If the interconnection node is used to connect a (set of) sub-ring(s) to another network, then there is no Ethernet ring accessed by two ring ports. Per ITU-T Rec.G.8032/Y.1344 (03/2010).

Interface

In IPv6, a node's attachment to a link.

Interface identifier

In IPv6, a link-dependent identifier for an interface that is (at least) unique per link. Stateless address autoconfiguration combines an interface identifier with a prefix to form an address. In address autoconfiguration, an interface identifier is a bit string of known length. The exact length of an interface identifier and the way it is created is defined in a separate link-type specific document that covers issues related to the transmission of IP over a particular link type. In many cases, the identifier will be the same as the interface's link- layer address.

Invalid address

In IPv6, an address that is not assigned to any interface. A valid address becomes invalid when its valid lifetime expires. Invalid addresses should not appear as the destination or source address of a packet. In the former case, the internet routing system will be unable to deliver the packet, in the later case the recipient of the packet will be unable to respond to it.

IP

IP (Internet Protocol) is a protocol used for communicating data across an internet network. IP is a "best effort" system, which means that no packet of information sent over is assured to reach its destination in the same condition it was sent. Each device connected to a Local Area Network (LAN) or Wide Area Network (WAN) is given an Internet Protocol address, and this IP address is used to identify the device uniquely among all other devices connected to the extended network.

The current version of the Internet protocol is IPv4, which has 32-bits Internet Protocol addresses allowing for in excess of four billion unique addresses. This number is reduced drastically by the practice of webmasters taking addresses in large blocks, the bulk of which remain unused. There is a rather substantial movement to adopt a new version of the Internet Protocol, IPv6, which would have 128-bits Internet Protocol addresses. This number can be represented roughly by a three with thirty-nine zeroes after it. However, IPv4 is still the protocol of choice for most of the Internet.

The IPv6 header fields are:

Version: The 4-bit Internet Protocol version number (6).

Traffic Class: An 8-bit traffic class field.

Flow Label: A 20-bit flow label.

Payload Length: the 16-bit unsigned integer. The Length of the IPv6 payload (i.e., the rest of the packet following this IPv6 header, in octets. Note that any extension headers present are considered part of the payload (i.e., included in the length count).

Next Header: An 8-bit selector that identifies the type of header immediately following the IPv6 header. Uses the same values as the IPv4 Protocol field.

Hop Limit: An 8-bit unsigned integer decremented by 1 by each node that forwards the packet. The packet is discarded if the Hop Limit is decremented to zero.

Source Address: The 128-bit address of the originator of the packet.

Destination Address: The 128-bit address of the intended recipient of the packet (possibly not the ultimate recipient, if a Routing header is present).

A full IPv6 implementation also includes these six extension headers: Hop-by-Hop Options, Routing (Type 0), Fragment, Destination Options, Authentication, and Encapsulating Security Payload headers. Unlike IPv4, IPv6 nodes are not required to enforce a maximum packet lifetime, which is why the IPv4 "Time to Live" field was renamed "Hop Limit" in IPv6.

IPMC

IPMC (IP MultiCast) provides a means to talk to a group of hosts (a multicast group), where each host has a different MAC address, and at the same time ensure that other hosts, which are not part of the multicast group, don't process the information. Broadcast packets make use of a broadcast MAC address (FF:FF:FF:FF:FF:FF), which includes setting the broadcast/multicast bit in the address. (Unicast packets are delivered to a specific recipient on an Ethernet or IEEE 802.3 subnet by setting a specific layer 2 MAC address on the Ethernet packet address.) A multicast address is associated with a group of interested receivers. In IPv4, addresses 224.0.0.0 through 239.255.255.255 (the former Class D addresses) are designated as multicast addresses.[3] IPv6 uses the address block with the prefix ff00::/8 for multicast applications. In either case, the sender sends a single datagram from its unicast address to the multicast group address and the intermediary routers take care of making copies and sending them to all receivers that have joined the corresponding multicast group.

IP multicast is a method of sending Internet Protocol (IP) datagrams to a group of interested receivers in a single transmission. It is often employed for streaming media applications on the Internet and private networks. The method is the IP-specific version of the general concept of multicast networking. It uses special reserved multicast address blocks in IPv4 and IPv6. In IPv6, IP multicast addressing replaces broadcast addressing as implemented in IPv4. The Linux commands *ping* and *netstat* are helpful when using IP multicast.

Ping commands can be used for multicast addresses by providing a multicast address as argument. Running *netstat* with the *-g* option on a Linux system displays the set of all multicast groups that the Linux system has joined.

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.

J

Jumbo frames

The LIB-44xx supports jumbo frames. This frame size is set to 9600 bytes or jumbo mode by default. The frame size is configurable to any value from 1500-9600 bytes. The jumbo mode is applicable only for data plane traffic; management and control plane traffic is still restricted to 1500 bytes.

L

LACP

LACP is an IEEE 802.3ad standard protocol. The **L**ink **A**ggregation **C**ontrol **P**rotocol allows bundling several physical ports together to form a single logical port. LACP is used by neighbouring devices to agree on adding links to a Link Aggregation Group, and to maintain packet ordering within each LAG. LACP will form an aggregation when 2 or more LIB-44xx ports are connected to the same partner.

LCI

Location Configuration Information.

Link

In IPv6, a communication facility or medium over which nodes can communicate at the link layer (i.e., the layer immediately below IPv6). Examples are Ethernets (simple or bridged); PPP links; X.25, Frame Relay, or ATM networks; and internet (or higher) layer "tunnels", such as tunnels over IPv4 or IPv6 itself.

Link-layer address

In IPv6, a link-layer identifier for an interface. Examples include IEEE 802 addresses for Ethernet links and E.164 addresses for ISDN links.

Link-local Address

One of IPv6 addresses for local link usage. In IPv6, an address having link-only scope that can be used to reach neighbouring nodes attached to the same link. All interfaces have a link-local unicast address.

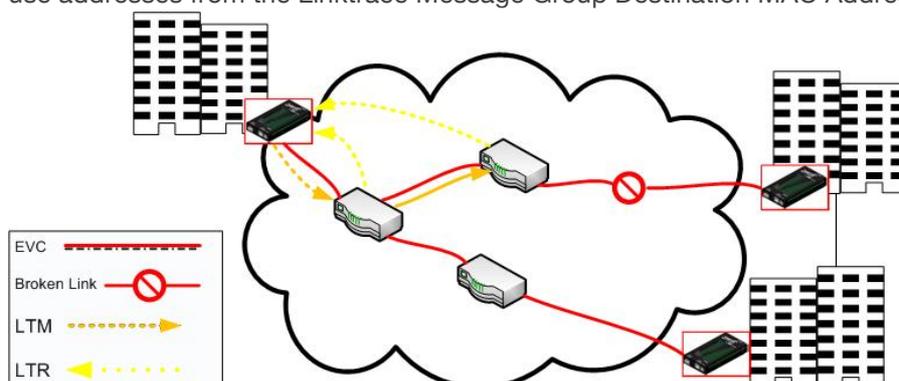
Link MTU

The IPv6 Maximum Transmission Unit - the maximum packet size in octets that can be conveyed over a link.

Link Trace

Link Trace messages are Multicast frames that a MEP transmits to track the path (hop-by-hop) to a destination MEP. Each receiving MEP sends a Linktrace Reply (LTR) directly to the Originating MEP, and regenerates the Linktrace Message: Each Linktrace Message (LTM) is a CFM PDU initiated by a MEP to trace a path to a target MAC address, forwarded from MIP to MIP, up to the point at which the LTM reaches its target, a MEP, or can no longer be forwarded. Each MP along the path to the target generates an LTR.

Each Linktrace Reply (LTR) is a unicast CFM PDU sent by an MP to a MEP, in response to receiving an LTM from that MEP. Linktrace Replies (LTRs) are carried in unicast frames. Linktrace Messages (LTMs) use addresses from the Linktrace Message Group Destination MAC Addresses table.



An LTM is used to signal to the MEP to transmit an LTM and to create an LTM entry in the MEP's Linktrace Database. The MA End Point can then be examined to determine whether or not the corresponding LTRs have been received by the MEP.

ETH-LT (Ethernet Link Trace) is an on-demand OAM function that can be used 1) to retrieve adjacency relationship between a MEP and a remote MEP or MIP, and 2) for Fault localization – when a fault (e.g., a link and/or a device failure) occurs, the sequence of MIPs and/or MEP will likely differ from the expected sequence. These differences provide information about the fault location.

ETH-LT request information is initiated in a MEP on an on-demand basis. After transmitting a frame with ETH-LT request information, the MEP expects to receive frames with ETH-LT reply information within a specified period of time. Network elements containing MIPs or MEPs and receiving the frame with ETH-LT request information respond selectively with frames containing ETH-LT reply information.

LLAG

(Local Link Aggregation Group) is one of two supported types of Link Aggregation Groups (same as Link Aggregation Group). With LLAG, all ports in an LLAG must reside on the same unit, any number of LLAGs may be configured for each unit in a stack, and each LLAG may consist of up to 16 ports. LLAGs are configured the same way as link aggregation groups for a standalone device (e.g., LIB-44xx). For both LLAGs and GLAGs, the egress port is chosen based on an 'aggregation code' that is calculated for the frame. This ensures that frames relating to a given frame flow are forwarded on the LLAG or GLAG member port, and thus do not risk being re-ordered. See also "GLAG".

GLAG

(Global Link Aggregation Group) is one of two supported types of Link Aggregation Groups. With GLAG, ports in a GLAG may reside on the same unit, up to two GLAGs are supported per stack, and each of the two GLAGs may consist of up to eight ports.

For both LLAGs and GLAGs, the egress port is chosen based on an 'aggregation code' that is calculated for the frame. This ensures that frames relating to a given frame flow are forwarded on the LLAG or GLAG member port, and thus do not risk being re-ordered. See also "LLAG".

LLC

The IEEE 802.2 Logical Link Control (LLC) protocol provides a link mechanism for upper layer protocols. It is the upper sub-layer of the Data Link Layer and provides multiplexing mechanisms that make it possible for several network protocols (IP, IPX) to coexist within a multipoint network. LLC header consists of 1 byte DSAP (Destination Service Access Point), 1 byte SSAP (Source Service Access Point), 1 or 2 bytes Control field followed by LLC information.

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).

LLDP-MED

LLDP-MED Link Layer Discovery Protocol Media Endpoint Discovery. LLDP-MED is an extension of IEEE 802.1ab and is defined by the Telecommunication Industry Association (TIA-1057).

LLDPDU

Link Layer Discovery Protocol Data Unit, as defined in IEEE 802.1AB.

LOAM

(Link OAM) Ethernet Connectivity Fault Management (CFM) provided per IEEE 802.3ah OAM. The major features covered by this protocol are Discovery, Link Monitoring, Remote Fault Detection, and Remote Loopback. The S3240 NIDs support both Link layer OAM (LOAM, per IEEE 802.3–2005

Clause 57) and Service layer OAM (SOAM, per IEEE 802.1AG and Y.1731). Compare to “SOAM”.

LOC

LOC (Loss Of Connectivity) is detected by a [MEP](#) and indicates lost connectivity in the network. LOC can be used as a switch criteria by [EPS](#).

M

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 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.

Major ring

The Ethernet ring that is connected on two ports to an interconnection node. From ITU-T Rec.G.8032/Y.1344 (03/2010).

Master clock

In the context of a single PTP communication path, a clock that is the source of time to which all other clocks on that path synchronize.

A system of 1588 clocks may be segmented into regions separated by boundary clocks. Within each region there will be a single clock, the master clock, serving as the primary source of time. These master clocks will in turn synchronize to other master clocks and ultimately to the grandmaster clock.

MD5

MD5 (**M**essage-**D**igest algorithm **5**) is a message digest algorithm used in a cryptographic hash function with a 128-bit hash value. It was designed by Ron Rivest in 1991. MD5 is officially defined in RFC 1321 - The MD5 Message-Digest Algorithm.

MD5 is an authentication protocol; one of two cryptography methods used for LIB-44xx user authentication. MD5 is a widely used cryptographic hash function with a 128-bit hash value. Specified in RFC 1321, MD5 is used in a wide range of security applications, and is also commonly used to check file integrity. However, it has been shown that MD5 is not collision resistant; as such, MD5 is not suitable for applications like SSL certificates or digital signatures that rely on this property. MD5 was designed by Ron Rivest in 1991 to replace the earlier hash function MD4. See also “SHA”.

ME

(Maintenance Entity) An entity that requires management and is a relationship between two maintenance entity group (MEG) end points. MEs in Ethernet networks can nest but not overlap.

MED

Media Endpoint Discovery.

MEG

(Maintenance Entity Group) A ME Group (MEG) consists of the MEs that belong to the same service inside a common OAM domain.

MEG Level

The MEG Level is used to distinguish between OAM frames belonging to different nested MEs. MEs belonging to the same MEG share a common MEG Level. Eight MEG Levels have been identified for the purposes of Ethernet OAM.

When a Subscriber, Service Providers, and Network Operators share the MEG Levels space, allocation of MEG Levels can be negotiated between the various roles involved. A default allocation of MEG Levels is such that Service OAM frames for a Subscriber ME use MEG Level 7, 6 or 5; Service OAM frames for

an EVC ME use MEG Level 3 or 4 as EVC ME belongs to a Service Provider OAM Domain; and Operator MEGs use MEG Levels 2, 1, or 0. The MEG Levels used for UNI ME and NNI ME default to 0. Note that this default allocation of MEG Level space between Subscribers, Service Providers and Operators could change based on a mutual agreement between them.

MEG level of a MEP (0-7). The defaults per MEF 30 are:

MEG	Default MEG Level	Suggested Use (MEF 30)
Subscriber MEG	6	Subscriber monitoring of an Ethernet service.
Test MEG	5	SP isolation of subscriber reported problems.
EVC MEG	4	SP monitoring of provided service.
Service Provider MEG	3	SP Monitoring of Service Provider network.
Operator MEG	2	Netw. Operator monitoring of the portion of a network.
UNI MEG	1	Service Provider monitoring of a UNI.
ENNI MEG	1	Network Operators' monitoring of an ENNI.

(where SP = Service Provider)

Note: Assignment of numerical MEG Levels to 'subscriber' (or customer) role, Service Provider role, and Operator role is somewhat arbitrary since those terms imply business relationships that cannot be standardized. For example, a 'subscriber' (or customer) may also be an Operator seeking a service from another Operator. The MEG Level default values are consistent with a shared MEG Level model across Subscriber, Operators, and Service Providers.

Note: The MEF and Broadband Forum (BBF) are not aligned on the use of MEG Level 5. If interworking between an MEF compliant implementation and a BBF compliant implementation is required, an agreement on the use of MEG Level 5 is required between the two parties.

MEP

A MEP (Maintenance Entity Endpoint) is an endpoint in a Maintenance Entity Group (ITU-T Y.1731). A MEP (Maintenance end point) is an inward-facing point at the edge of the domain that defines the boundary and confines CFM messages within these boundaries. Inward facing means that they communicate through the relay function side, not the wire side (connected to the port). See also MIP, Down MEP, and Up MEP.

A MEG End Point (MEP) is a provisioned OAM reference point which can initiate and terminate proactive OAM frames. A MEP can also initiate and react to diagnostic OAM frames. A Point-to-Point EVC has two MEPs, one on each end point of the ME. A Multipoint-to-Multipoint EVC of n UNIs has n MEPs, one on each end point.

MIP

(Maintenance intermediate point) – A point internal to a domain, not at the boundary, that responds to CFM only when triggered by trace route and loopback messages. MIPs forward CFM frames received from MEPs and other MIPs, drop all CFM frames at a lower level, and forward all CFM frames at a higher level, regardless of whether they are received from the relay or wire side.

A MEG Intermediate Point (MIP) is a provisioned OAM reference point that can react to diagnostic OAM frames initiated by MEPs. A MIP does not initiate proactive or diagnostic OAM frames. See also “MEP”.

Mirroring

For debugging network problems or monitoring network traffic, the LIB-44xx can be configured to mirror frames from multiple ports to a mirror port. (In this context, [mirroring](#) a frame is the same as copying the frame.)

Both incoming (source) and outgoing (destination) frames can be mirrored to the mirror port.

MLD

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. MLD is a component of the Internet Protocol Version 6 (IPv6) suite, and is used by IPv6 routers to discover multicast listeners on a directly attached link (much as IGMP is used in IPv4). MLD is embedded in ICMPv6 instead of using a separate protocol. MLDv1 is similar to IGMPv2 and MLDv2 similar to IGMPv3. The MLD protocol is described in RFC 3810 which

was updated by RFC 4604. Windows Vista and later support MLDv2. FreeBSD 8 includes support for MLDv2. The Linux kernel has supported MLDv2 since v 2.5.68.

MLD snooping

With MLD snooping, IPv6 multicast data is selectively forwarded to a list of ports that want to receive the data, instead of being flooded to all ports in a VLAN. This list of ports is created by 'snooping' IPv6 multicast control packets. In IPv6, MLD snooping performs a similar function to the IGMP snooping used in IPv4.

MSTI

An MSTI (**M**ultiple **S**panning **T**ree **I**nstance) is typically one of the uplink ports that connects to one of the gateway devices. Valid MSTI ID values are from 0 through 4094. MSTI information can include VLAN mapping, bridge priority, port priority, and cost. MSTP allows formation of MST regions that can run multiple MST instances (MSTI). Multiple regions and other STP bridges are interconnected using one single common spanning tree (CST). Unlike some proprietary per-VLAN spanning tree implementations, MSTP includes all of its spanning tree information in a single BPDU format. This reduces the number of BPDUs required on a LAN to communicate spanning tree information for each VLAN, and also ensures backward compatibility with RSTP (and effectively, classic STP too). MSTP does this by encoding additional region information after the standard RSTP BPDU as well as a number of MSTI messages (0 to 64 instances; many bridges support fewer). Each of these MSTI configuration messages conveys the spanning tree information for each instance. Each instance can be assigned a number of configured VLANs, and frames (packets) assigned to these VLANs operate in this spanning tree instance whenever they are inside the MST region. To avoid conveying their entire VLAN to spanning tree mapping in each BPDU, bridges encode an MD5 digest of their VLAN-to-instance table in the MSTP BPDU. This digest is then used by other MSTP bridges, along with other administratively configured values, to determine if the neighbouring bridge is in its MST region. See also "CIST".

Multicast address

In IPv6, an identifier for a set of interfaces (typically belonging to different nodes). A packet sent to a multicast address is delivered to all interfaces identified by that address.

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 (Wikipedia).

N

NAS

The NAS (**N**etwork **A**ccess **S**erver) 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](#).

NDP

(Neighbour Discovery Protocol) - a protocol in the Internet Protocol Suite used with IPv6. NDP operates in the Link Layer and is responsible for address autoconfiguration of nodes, discovery of other nodes on the link, determining the Link Layer addresses of other nodes, duplicate address detection, finding available routers and DNS servers, address prefix discovery, and maintaining reachability information about the paths to other active neighbour nodes per IETF RFC 4861.

Neighbours

Nodes attached to the same link. Per IETF RFC 2461, Neighbour Discovery for IPv6 is done by Sending Router Advertisements and processing Router Solicitation.

NENA

National Emergency Number Association, the body responsible for evolution of public ECS architectures in North America.

NetBIOS

NetBIOS (Network Basic Input/Output System) is a program that allows applications on separate computers to communicate within a Local Area Network (LAN), and it is not supported on a Wide Area Network (WAN). The NetBIOS giving each computer in the network both a NetBIOS name and an [IP](#) address corresponding to a different host name, provides the session and transport services described in the Open Systems Interconnection (OSI) model.

NFS

NFS (Network File System) allows hosts to mount partitions on a remote system and use them as though they are local file systems. NFS allows the system administrator to store resources in a central location on the network, providing authorized users continuous access to them, which means NFS supports sharing of files, printers, and other resources as persistent storage over a computer network.

NID

Network Interface Devices - a NID is an effective way of providing operational and capital savings to service providers. A NID installs at the customer premise and provides a demarcation point between the service provider and customer's network. NIDs allow for end-to-end Operations, Administration and Maintenance (OAM) functionality for the service provider. The basic functions, such as loopback testing and remote fault isolation, in the NID provide service providers a number of benefits, including: reduced truck rolls, fewer test sets in the field, and increased reliability. The result of this for the service provider is a reduction in OpEx and CapEx while providing a faster return on investment (ROI).

While the operational savings of NIDs can be shown with their features and capabilities for remote troubleshooting, easy installation and SLA monitoring to reduce SLA penalties, it is important for service providers to be aware of the additional revenue streams and services that can be achieved when using a NID at the demarcation point. NIDs have advanced features such as bandwidth allocation, QoS, VLAN and other features that allow the service provider the capability to provide tiered service offerings to customers.

NMS

Network Management System. Contrast "EMS".

NNI

(Network to Network Interface) In carrier Ethernet, the demarcation / peering point between service providers (ENNI) or between service provider internal networks (I-NNI), per MEF 3, 12, 17, 4, 30, 31.

Node

In IPv6, a device that implements IPv6.

Non-revertive mode

In non-revertive mode of unidirectional protection switching operation, in conditions where working traffic is being transmitted via the protection entity, if local protection switching requests have been previously active and now become inactive, a local "do-not-revert state" is entered. Since this state now represents the highest priority local request, it is indicated on the transmitted "request/state" information and maintains the switch, preventing reversion back to the released bridge-selector position in non-revertive mode under no-request conditions. With bidirectional protection switching operation, a local do-not-revert state is entered when there is no higher priority of request received from the far end than that of the do-not-revert state, or when both the local state and far-end state are NR with the requested signal number 1.

Generally, Revertive operation is useful when the working transport entity is more optimized or the protection transport entity carries best effort traffic; Non-revertive operation can minimize the number of switching and service outage time. See also "Revertive mode".

NTP

NTP (Network Time Protocol) is a network protocol for synchronizing the clocks of computer systems. NTP uses [UDP](#) (datagrams) as transport layer.

O**OAM**

OAM (Operation Administration and Maintenance) protocol is described in ITU-T Y.1731 and is used to implement carrier ethernet functionality. [MEP](#) functionality like [CC](#) and [RDI](#) is based on this. The LIB-44xx provides configuration and monitoring of two types of Ethernet OAM:

- 1) end-to-end service OAM (SOAM) per IEEE 802.1ag and ITU-T Y.1731, to let Ethernet service providers monitor their services proactively, measure end-to-end performance, and guarantee that the customers receive the contracted SLA. Fault monitoring and performance measurement include frame delay, frame delay variation, frame loss and availability.
- 2) single segment Link OAM (LOAM) per IEEE 802.3ah for remote management and fault indication, including remote loopback, dying gasp, and MIB parameter retrieval.

OID

(Object Identifier) Many standards define certain objects that require unambiguous identification, which can be achieved by 'registration'. Registration is the assignment of an object identifier (OID) to an object in a way which makes the assignment available to interested parties. It is carried out by a registration authority. Registration can be effected by publishing in the standard the names and the corresponding definitions of object. Such a mechanism requires amendment of the standard for each registration, and hence is not appropriate in cases where the registration activity is high. Alternatively, registration can be affected by letting organizations act as registration authorities to perform registration on a flexible basis. The registration tree is managed in a completely decentralized way (a node gives full power to its children) and it is impossible to be exhaustive (particularly world-wide). The registration tree is defined and managed following the ITU-T X.660 & X.670 Recommendation series (or the ISO/IEC 9834 series of International Standards).

One-step clock

A clock that provides time information using a single event message.

Optional TLVs

An LLDP frame contains multiple [TLVs](#). For some TLVs it is configurable if the LIB-44xx includes 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.

Ordinary clock

A clock that has a single Precision Time Protocol port in a domain and maintains the timescale used in the domain. It may serve as a source of time (i.e., be a master clock), or may synchronize to another clock (i.e., be a slave clock). An ordinary clock is a 1588 clock with a single PTP port.

OUI

(Organizationally Unique Identifier) 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.

P**Packet**

An IPv6 header plus payload.

Parent clock

The master clock to which a clock is synchronized.

Path MTU

The minimum IPv6 link MTU of all the links in a path between a source node and a destination node.

Path Cost

A path cost value is given to each port. The cost is usually based on 802.1d guidelines. According to the original specification, cost is 1,000 Mbps (1 gigabit per second) divided by the bandwidth of the segment connected to the port. Therefore, a 10 Mbps connection would have a cost of (1,000/10) 100. A 'path cost' is an administratively assigned value for the contribution of a port to the path cost of paths toward the spanning tree root. A value of '0' assigns the automatically calculated default Path Cost value to the port (the default Path Cost). This complements the object dot1dStpPortPathCost or dot1dStpPortPathCost32, which returns the operational value of the path cost.

Typical STP path costs are shown below for certain data rates.

<u>Data rate</u>	<u>STP Cost (802.1D-1998)</u>	<u>RSTP Cost (802.1W-2001)</u>
4 Mbps	250	5,000,000
10 Mbps	100	2,000,000
16 Mbps	62	1,250,000
100 Mbps	19	200,000
1 Gbps	4	20,000
2 Gbps	3	10,000
10 Gbps	2	2,000

The recommended values for any intermediate link speed can be calculated as 20,000,000,000/(Link Speed in Kb/s). Limiting the range of the Path Cost parameter to 1-200,000,000 ensures that the accumulated Path Cost cannot exceed 32 bits over a concatenation of 20 hops.

PCP

PCP (Priority Code Point) is a 3-bit field storing the priority level for the 802.1Q frame (also known as [User Priority](#)).

PD

PD (Powered Device) in a [PoE](#) system the power is delivered from a PSE (power sourcing equipment) to a remote device. The remote device is called a PD.

PDU

(Protocol Data Units) 1. Information that is delivered as a unit among peer entities of a network and that may contain control information, address information or data. 2. In a layered system, a unit of data which is specified in a protocol of a given layer and which consists of protocol control information and possibly user data of that layer.

Peer-to-peer transparent clock

A transparent clock that, in addition to providing Precision Time Protocol event transit time information, also provides corrections for the propagation delay of the link connected to the port receiving the PTP event message. In the presence of peer-to-peer transparent clocks, delay measurements between slave clocks and the master clock are performed using the peer-to-peer delay measurement mechanism.

PerfectReach™

One of the LIB-44xx energy efficient modes. An intelligent algorithm that actively determines the needed power level based on cable length.

Per flow QoS

The ability to identify a traffic flow, enable rules on how that specific flow should be treated, and then define how the flow should behave when forwarded with other traffic flows. See also "flow".

PHY

PHY (Physical Interface Transceiver) is the device that implements the Ethernet physical layer per IEEE-802.3.

PIC

(Peripheral interface controller) a family of specialized microcontroller chips.

PING

The ping program sends a series of packets over a network or the Internet to a specific computer in order to generate a response from that computer. The other computer responds with an acknowledgment that it received the packets. Ping was created to verify whether a specific computer on a network or the Internet exists and is connected. `Ping` uses Internet Control Message Protocol ([ICMP](#)) packets. The PING Request is the packet from the origin computer, and the PING Reply is the packet response from the target.

PoE

PoE (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.

Policer

A policer can limit the bandwidth of received frames. It is located in front of the ingress queue. Ethernet frame color is determined using a bandwidth profile through a traffic management function called "policing".

Color	Conformance	Ethernet Frame Delivery Expectation
Green	Conformant to CIR	Frames coloured green and delivered per the SLO.
Yellow	Non-conformant to CIR Conformant to EIR	Frames coloured yellow and may be delivered but with no SLO assurances.
Red	Non-conformant to CIR or EIR	Frames coloured red and dropped

POP3

POP3 (Post Office Protocol version 3) is a protocol for email clients to retrieve email messages from a mail server. POP3 is designed to delete mail on the server as soon as the user has downloaded it. However, some implementations allow users or an administrator to specify that mail be saved for some period of time. POP can be thought of as a "store-and-forward" service.

An alternative protocol is Internet Message Access Protocol ([IMAP](#)). IMAP provides the user with more capabilities for retaining e-mail on the server and for organizing it in folders on the server. IMAP can be thought of as a remote file server. POP and IMAP deal with the receiving of e-mail and are not to be confused with the Simple Mail Transfer Protocol ([SMTP](#)). You send e-mail with SMTP, and a mail handler receives it on your recipient's behalf. Then the mail is read using POP or IMAP. IMAP4 and POP3 are the two most prevalent Internet standard protocols for e-mail retrieval. Virtually all modern e-mail clients and servers support both.

PPPoE

PPPoE (Point-to-Point Protocol over Ethernet) network protocol is used for encapsulating Point-to-Point Protocol (PPP) frames inside Ethernet frames. It is used mainly with ADSL services where individual users connect to the ADSL transceiver (modem) over Ethernet and in plain Metro Ethernet networks (Wikipedia).

Preferred address

In IPv6, an address assigned to an interface whose use by upper layer protocols is unrestricted. Preferred addresses may be used as the source (or destination) address of packets sent from (or to) the interface.

Preferred lifetime

In IPv6, the length of time that a valid address is preferred (i.e., the time until deprecation). When the preferred lifetime expires, the address becomes deprecated.

Private VLAN

In a private VLAN, communication between ports in that private [VLAN](#) is not permitted. A VLAN can be configured as a private VLAN (PVLAN).

PTP

PTP (Precision Time Protocol) is a network protocol for synchronizing the clocks of computer systems.

Q

QCE

QCE (QoS Control Entry) describes [QoS](#) class associated with a particular QCE ID. There are six QCE frame types: [Ethernet Type](#), [VLAN](#), [UDP/TCP Port](#), [DSCP](#), [TOS](#), and [Tag Priority](#). Frames can be classified by one of four different QoS classes: "Low", "Normal", "Medium", and "High" for individual application.

QCL

QCL (QoS Control List) is the list table of [QCEs](#), containing [QoS](#) control entries that classify to a specific QoS class on specific traffic objects. Each accessible traffic object contains an identifier to its QCL. The privileges determine specific traffic object to specific QoS class.

QL

QL, in [SyncE](#), is the Quality Level of a given clock source. This is received on a port in a [SSM](#) indicating the quality of the clock received in the port.

QoS

QoS (Quality of Service) 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 secret to a successful end-to-end business solution. QoS is the set of techniques to manage network resources.

When discussing QoS features:

- "Packets" carry traffic at Layer 3.
- "Frames" carry traffic at Layer 2 (Layer 2 frames carry Layer 3 packets).
- "Classification" is the selection of traffic to be marked.
- "Marking" (per RFC 2475) is the process of setting a Layer 3 DSCP value of a packet.
- "Policing" is limiting bandwidth used by a flow of traffic; policing can either mark or drop traffic.

R

R-APS

R-APS is an acronym for Ring APS. Per G.8032v1, in ERPS there is a central node called the 'RPL Owner Node' which blocks one of the ports to ensure that there is no loop formed for the Ethernet traffic. The link blocked by the RPL Owner node is called the Ring Protection Link or RPL. The node at the other end of the RPL is known as RPL Neighbour Node. It uses R-APS control messages to coordinate the activities of switching on/off the RPL link.

Any failure along the ring triggers an R-APS(SF) (R-APS signal fail) message along both directions from the nodes adjacent to the failed link after these nodes have blocked the port facing the failed link. On obtaining this message, RPL owner unblocks the RPL port. Note that a single link failure anywhere in the ring ensures a loop free topology.

During the recovery phase when the failed link gets restored, the nodes adjacent to the restored link send R-APS (NR) (R-APS no request) messages. On obtaining this message, the RPL owner blocks the RPL port and then sends a R-APS (NR, RB) (R-APS no request, root blocked) messages. This causes all other nodes other than the RPL Owner in the ring to unblock all of the blocked ports.

This protocol is robust enough to work for unidirectional failure and multiple link failure scenarios in a ring topology. It allows mechanism to force switch (FS) or manual switch (MS) to support field maintenance scenarios.

R-APS virtual channel

The Ring Automatic Protection Switching (R-APS) channel connection between two interconnection nodes of a sub-ring in (an)other Ethernet ring(s) or network(s). Its connection characteristics (e.g., path, performance, etc.) are influenced by the characteristics of the network (e.g., Ethernet ring) providing connectivity between the interconnection nodes. From ITU-T Rec.G.8032/Y.1344 (03/2010).

RARP

RARP (**R**everse **A**ddress **R**esolution **P**rotocol) is a protocol used to obtain an [IP](#) address for a given hardware address, such as an Ethernet address. RARP is the complement of [ARP](#).

RADIUS

RADIUS (**R**emote **A**uthentication **D**ial **I**n **U**ser **S**ervice) is a networking protocol that provides centralized access, authorization and accounting management for people or computers to connect and use a network service.

RDI

RDI (Remote Defect Indication) is an [OAM](#) function used by a [MEP](#) to indicate a defect detected to the remote peer MEP. The IEEE Remote Defect Indication (RDI) is a single bit carried by the CCM. The absence of RDI in a CCM indicates that the transmitting MEP is receiving CCMs from all configured MEPs.

A MEP can use ITU-T ETH-RDI to notify its peer MEPs that it detects a defect condition. ETH-RDI is used only if ETH-CC transmission is enabled. ETH-RDI is used in single-ended fault management and in contributing to far-end performance monitoring. A MEP in a defect condition transmits frames with ETH-RDI information. When a MEP receives frames with ETH-RDI information it determines that its peer MEP has encountered a defect condition

A MEP, on detecting a defect condition with its peer MEP, sets the RDI field in the CCM frames for the duration of the defect condition. CCM frames are transmitted periodically based on the CCM transmission period when the MEP is enabled for CCM frames transmission. When the defect condition clears, the MEP clears the RDI field in the CCM frames in subsequent transmissions.

Reversion Time (WTR time)

In revertive mode, the Reversion Time is the difference between the repair instant of the original resource and the Reversion Instant.

Revertive Mode

Protection is in revertive mode if, after a resource failure and its subsequent repair, the network automatically reverts to using this initial resource. The protection is in non-revertive mode otherwise. Automatic reversion may include a reversion timer (i.e., the Wait To Restore), which delays the time of reversion after the repair.

In Revertive mode of unidirectional protection switching operation, in conditions where working traffic is being received via the protection entity, if local protection switching requests have been previously active and now become inactive, a local wait-to-restore state is entered. Since this state now represents the highest priority local request, it is indicated on the transmitted "request/state" information and maintains the switch. With bidirectional protection switching, a local wait-to-restore state is entered only when there is no higher priority of request received from the far end than that of the wait-to-restore state. This state normally times out and becomes a no request state after the wait-to-restore timer has expired. The wait-to-restore timer is deactivated earlier if any local request of higher priority pre-empts this state. A switch to the protection entity may be maintained by a local wait-to-restore state or by a remote request (wait-to-restore or other) received via the "request/state" information. So, in a case where a bidirectional failure for a working entity has occurred and subsequent repair has taken place, the bidirectional reversion back to the working entity does not take place until both wait-to-restore timers at both ends have expired. See also "Non-revertive mode".

Ring MEL

The Maintenance Entity Group (MEG) Level providing a communication channel for ring automatic protection switching (R-APS) information. From ITU-T Rec.G.8032/Y.1344 (03/2010).

Ring Protection Link (RPL)

The ring link that under normal conditions, i.e., without any failure or request, is blocked (at one or both ends) for traffic channel, to prevent the formation of loops. From ITU-T Rec.G.8032/Y.1344 (03/2010).

RPL Neighbour node

The RPL neighbour node, when configured, is an Ethernet ring node adjacent to the RPL that is responsible for blocking its end of the RPL under normal conditions (i.e., the ring is established and no requests are present in the ring) in addition to the block by the RPL owner node. However, it is not responsible for activating the reversion behaviour. From ITU-T Rec.G.8032/Y.1344 (03/2010). Contrast "RPL Owner Node".

RPL Owner node

The RPL owner node is an Ethernet ring node adjacent to the RPL that is responsible for blocking its end of the RPL under normal conditions (i.e., the ring is established and no requests are present in the ring). Furthermore, it is responsible for activating reversion behaviour from protected or manual switch/forced switch (MS/FS) conditions. From ITU-T Rec.G.8032/Y.1344 (03/2010). Contrast "RPL Neighbour Node".

Router

In IPv6, a node that forwards IPv6 packets not explicitly addressed to itself.

Router Port

A port on the Ethernet switch that leads switch towards the Layer 3 multicast device.

RSTP

In 1998, the IEEE with document 802.1w introduced an evolution of [STP](#): the **R**apid **S**panning **T**ree **P**rotocol, which provides for faster spanning tree convergence after a topology change. Standard IEEE 802.1D-2004 now incorporates RSTP and obsoletes STP, while at the same time being backwards-compatible with STP.

S**SA**

(Source Address); contrast "DA."

SAMBA

Samba is a program running under UNIX-like operating systems that provides seamless integration between UNIX and Microsoft Windows machines. Samba acts as file and print servers for Microsoft Windows, IBM OS/2, and other SMB client machines. Samba uses the Server Message Block (SMB) protocol and Common Internet File System (CIFS), which is the underlying protocol used in Microsoft Windows networking. Samba can be installed on a variety of operating system platforms, including Linux, most common Unix platforms, OpenVMS, and IBM OS/2. Samba can also register itself with the master browser on the network so that it would appear in the listing of hosts in Microsoft Windows "Neighbourhood Network".

SFP

Small Form Factor Pluggable module (1-4Gbps).

SFP+

Small Form Factor Pluggable Plus module (8-10Gbps).

SHA

SHA (**S**ecure **H**ash **A**lgorithm) is designed by the National Security Agency (NSA) and published by the NIST as a U.S. Federal Information Processing Standard. Hash algorithms compute a fixed-length digital representation (known as a message digest) of an input data sequence (the message) of any length. SHA is an authentication protocol; one of two cryptography methods used for LIB-44xx user

authentication. SHA-1 is a cryptographic hash function designed by the National Security Agency (NSA) and published by the NIST as a U.S. FIPS standard. SHA-1 is part of many widely accepted security applications and protocols (e.g., TLS, SSL, PGP, SSH, S/MIME, and IPSec). See also “MD5”.

Shaper

A shaper can limit the bandwidth of transmitted frames. It is located after the ingress queues.

Site-local Address

An IPv6 addresses for local site only. In IPv6, an address having scope that is limited to the local site.

SLO

A service level objective () is a key element of a service level agreement (SLA) between a service provider and a customer. SLOs are agreed as a means of measuring the performance of the Service Provider and are outlined as a way of avoiding disputes between the two parties based on misunderstanding.

The SLO may conclude one or more QoS measurements that are combined to produce the SLO achievement value.

SMAC

(Source MAC Address) The 12 hex digits of a source MAC address consist of the first/left 6 digits (which should match the vendor of the Ethernet NIC) and the last/right 6 digits which specify the interface serial number for that interface controller vendor. See also "DMAC". The list below identifies some of the blocks of assigned vendor MAC addresses (i.e., the first 3 bytes of a MAC source address).

00000C Cisco
 00000E Fujitsu
 00000F NeXT
 00001D Cabletron
 000022 Visual Technology
 00002A TRW
 00005A S & Koch
 00005E IANA
 000065 Network General
 00006B MIPS
 000093 Proteon
 08005A IBM
 080067 Comdesign
 080069 Silicon Graphics
 08007C Vitalink TransLAN III
 080080 XIOS
 080086 Imagen/QMS
 080087 Xyplex terminal servers
 080089 Kinetics AppleTalk-Ethernet interface
 080090 Retix Inc. Bridges

SMB

(SubMiniature version B) connectors are coaxial RF connectors developed in the 1960s. SMB connectors are smaller than SMA connectors. SMB connectors feature a snap-on coupling and are available in either 50 Ω or 75 Ω impedance. They offer excellent electrical performance from DC to 4 GHz. An SMB jack has a male centre pin, while an SMB plug has a female basket. Connectors are available for two SMB cable sizes; a Cable 2.6/50+75 S (3 mm outer / 1.7 mm inner diameter) and a Cable 2/50 S (2.2 mm outer / 1 mm inner diameter).

SMTP

SMTP (Simple Mail Transfer Protocol) is a text-based protocol that uses the Transmission Control Protocol ([TCP](#)) and provides a mail service modelled on the [FTP](#) file transfer service. SMTP transfers mail messages between systems and notifications regarding incoming mail.

SNAP

The SubNetwork Access Protocol (SNAP) is a mechanism for multiplexing, on networks using IEEE 802.2 LLC, more protocols than can be distinguished by the 8-bit 802.2 Service Access Point (SAP) fields. SNAP supports identifying protocols by Ethernet type field values; it also supports vendor-private protocol identifier.

SNMP

SNMP (Simple Network Management Protocol) is part of the Transmission Control Protocol/Internet Protocol ([TCP/IP](#)) protocol for network management. SNMP allow diverse network objects to participate in a network management architecture. It enables network management systems to learn network problems by receiving traps or change notices from network devices implementing SNMP.

SNTP

SNTP (Simple Network Time Protocol) is a network protocol for synchronizing the clocks of computer systems. SNTP uses [UDP](#) (datagrams) as transport layer.

SOAM

(Service OAM) provides Ethernet Connectivity Fault Management (CFM) per IEEE 802.1AG. Ethernet CFM comprises three protocols that work together to help administrators debug Ethernet networks: continuity check, link trace and loopback protocols. The LIB-44xx supports both Link layer OAM (LOAM, per IEEE 802.3–2005 Clause 57) and Service layer OAM (SOAM, per IEEE 802.1AG and Y.1731). Compare to 'LOAM'.

Solicited-node multicast address

In IPv6, a multicast address to which Neighbour Solicitation messages are sent. The algorithm for computing the address is given in Discovery.

Spanning Tree

The original spanning-tree protocol (STP) was created to prevent broadcast storms and other unwanted side effects of looping. Since, STP has been standardized as the 802.1d specification by the IEEE.

A spanning tree uses a spanning-tree algorithm (STA) to sense that the switch has more than one way to communicate with a node, then determine the best way, and then block out all other paths. The STA also keeps track of the other paths, in case the primary path becomes unavailable.

Each switch is assigned a group of IDs - one for the switch itself and one for each port on the switch. A switch's bridge ID (BID) is 8 bytes long and contains a bridge priority (2 bytes) along with one of the switch's MAC addresses (6 bytes). Each port ID is 16 bits long with two parts - a 6-bit priority setting and a 10-bit port number. A 'path cost' value is assigned to each port. See also "Path Cost".

SPROUT

Stack Protocol using Routing Technology is an advanced protocol for almost instantaneous discovery of topology changes within a stack as well as election of a master switch. SPROUT also calculates parameters for setting up each switch to perform shortest path forwarding within the stack.

SSID

Service Set Identifier is a name used to identify the particular 802.11 wireless LANs to which a user wants to attach. A client device will receive broadcast messages from all access points within range advertising their SSIDs, and can choose one to connect to based on pre-configuration, or by displaying a list of SSIDs in range and asking the user to select one (Wikipedia).

SSH

SSH (**S**ecure **S**hell) is a network protocol that allows data to be exchanged using a secure channel between two networked devices. The encryption used by SSH provides confidentiality and integrity of data over an insecure network. The goal of SSH was to replace the earlier rlogin, [TELNET](#) and rsh protocols, which did not provide strong authentication or guarantee confidentiality (Wikipedia).

SSM

SSM in [SyncE](#) is an abbreviation for Synchronization Status Message and contains a [QL](#) indication.

SSM

SSM (Source-Specific Multicast) IP version 4 (IPv4) addresses in the 232/8 (232.0.0.0 to 232.255.255.255) range are designated as source-specific multicast (SSM) destination addresses and

are reserved for use by source-specific applications and protocols. For IP version 6 (IPv6), the address prefix FF3x::/32 is reserved for source-specific multicast use. IETF RFC 4607 defines an extension to the Internet network service that applies to datagrams sent to SSM addresses and defines the host and router requirements to support this extension.

Source-specific multicast (SSM) is a method of delivering multicast packets in which the only packets that are delivered to a receiver are those originating from a specific source address requested by the receiver. By so limiting the source, SSM reduces demands on the network and improves security. SSM requires that the receiver specify the source address and explicitly excludes the use of the (*,G) join for all multicast groups in RFC 3376, which is possible only in IPv4's IGMPv3 and IPv6's MLDv2.

The burden of source discovery on the network can be significant with a large number of sources. In the SSM model, in addition to the receiver expressing interest in traffic to a multicast address, the receiver expresses interest in receiving traffic from only one specific source sending to that multicast address. This relieves the network of discovering many multicast sources and reduces the amount of multicast routing information that the network must maintain. SSM requires support in last-hop routers and in the receiver's operating system. SSM support is not required in other network components, including routers and even the sending host. Interest in multicast traffic from a specific source is conveyed from hosts to routers using IGMPv3 as specified in RFC 4607. SSM destination addresses must be in the range of 232.0.0.0/8 for IPv4 or FF3x::/96 for IPv6.

SSM identifies a set of multicast hosts not only by group address but also by source. An SSM group, called a 'channel', is identified as (S,G) where S is the source address and G is the group address.

Stateless auto-configuration

A process to get IPv6 addresses from IPv6 standards.

Stateless

A communications protocol that treats each request as an independent transaction that is unrelated to any previous request so that the communication consists of independent pairs of requests and responses.

A stateless protocol does not require the server to retain session information or status about each communications partner for the duration of multiple requests. An example of a stateless protocol is the Hypertext Transfer Protocol (HTTP) which is the foundation of data communication for the World Wide Web.

The stateless design simplifies the server design because there is no need to dynamically allocate storage to deal with conversations in progress. If a client dies in mid-transaction, no part of the system needs to be responsible for cleaning the present state of the server. A disadvantage of statelessness is that it may be necessary to include additional information in every request, and this extra information will need to be interpreted by the server. An example of a stateless protocol is HTTP. The protocol provides no means of storing a user's data between requests. As a work-around, HTTP Servers implement various session management methods, typically utilizing a unique identifier in a cookie or parameter that allows the server to track requests originating from the same client. Contrast this with a traditional FTP server that conducts an interactive session with the user. During the session, a user is provided a means to be authenticated and set various variables (working directory, transfer mode), all stored on the server as part of the user's state. From Wikipedia.

STP

Spanning **T**ree **P**rotocol is an OSI layer-2 protocol which ensures a loop free topology for any bridged LAN. The original STP protocol is now obsolete by [RSTP](#).

Sub-ring

An Ethernet ring which is connected to one or more other Ethernet rings or networks through using a pair of interconnection nodes. On their own, the sub-ring links do not form a closed loop. A closed connection of traffic may be formed by the sub-ring links and one or more links that are controlled by another Ethernet ring or network, between interconnection nodes. From ITU-T Rec.G.8032/Y.1344 (03/2010).

Sub-ring link

A span (e.g., link/port) connecting adjacent sub-ring nodes that is under the control of the Ethernet ring protocol control process (ERP control process) of the sub-ring. From ITU-T Rec.G.8032/Y.1344 (03/2010).

Subnet Mask (Address Mask)

A bit mask used to identify which bits in an IP address correspond to the network and subnet portions of the address. Referred to as the 'subnet' mask because the network portion of the address (the network mask) can be determined by the encoding inherent in an IP address.

SyncE

SyncE (Synchronous Ethernet) functionality is used to make a network 'clock frequency' synchronized. Not to be confused with real time clock synchronized (IEEE 1588).

Synchronized clocks

Two clocks are synchronized to a specified uncertainty when they have the same epoch and their measurements of the time of a single event at an arbitrary time differ by no more than that uncertainty.

T**TACACS+**

TACACS+ (**T**erminal **A**ccess **C**ontroller **A**ccess **C**ontrol **S**ystem **P**lus) 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.

Tag

An optional field in a frame header. In MEF 26 it is the 4-byte field that, when present in an Ethernet frame, appears immediately after the Source Address, or another tag in an Ethernet frame header and which consists of the 2-byte Tag Protocol Identification Field (TPID) which indicates S-Tag or C-Tag, and the 2-byte Tag Control Information field (TCI) which contains the 3-bit Priority Code Point, and the 12-bit VLAN ID field.

Tag Priority

Tag Priority is a 3-bit field storing the priority level for the 802.1Q frame.

TCP

TCP (Transmission Control Protocol) is a communications protocol that uses the Internet Protocol ([IP](#)) to exchange the messages between computers. The TCP protocol guarantees reliable and in-order delivery of data from sender to receiver and distinguishes data for multiple connections by concurrent applications (for example, Web server and e-mail server) running on the same host.

The applications on networked hosts can use TCP to create connections to one another. It is a connection-oriented protocol, which means that a connection is established and maintained until such time as the message or messages to be exchanged by the application programs at each end have been exchanged. TCP is responsible for ensuring that a message is divided into the packets that IP manages and for reassembling the packets back into the complete message at the other end. Common network applications that use TCP include the World Wide Web (WWW), e-mail, and File Transfer Protocol ([FTP](#)).

Telnet

TELNET (**TEL**etype **NET**work) is a terminal emulation protocol that uses the Transmission Control Protocol ([TCP](#)) and provides a virtual connection between TELNET server and TELNET client. TELNET enables the client to control the server and communicate with other servers on the network. To start a Telnet session, the client user must log in to a server by entering a valid username and password. Then, the client user can enter commands through the Telnet program just as if they were entering commands directly on the server console.

Tentative address

In IPv6, an address whose uniqueness on a link is being verified, prior to its assignment to an interface. A tentative address is not considered assigned to an interface in the usual sense. An interface discards received packets addressed to a tentative address, but accepts Neighbour Discovery packets related to Duplicate Address Detection for the tentative address.

TFTP

TFTP (**T**rivial **F**ile **T**ransfer **P**rotocol) is a transfer protocol that uses the User Datagram Protocol (**UDP**) and provides file writing and reading, but it does not provide directory service and security features.

Throttling

An LIB-44xx function used to limit the number of multicast groups to which a switch port can belong.

ToS

ToS (**T**ype **o**f **S**ervice) is implemented as the IPv4 ToS priority control. It is fully decoded to determine the priority from the 6-bit ToS field in the **IP** header. The most significant 6 bits of the ToS field are fully decoded into 64 possibilities, and the singular code that results is compared against the corresponding bit in the IPv4

ToS priority control bit (0~63).

TLV

TLV (**T**ype **L**ength **V**alue). An LLDP frame can contain multiple pieces of information. Each of these pieces of information is known as TLV. For the Type, Length, Value format, LLDP frames are sent by each equipment on each port at a fixed frequency. A frame contains a Link Layer Discovery Protocol Data Unit (LLDPDU) which is a set of type, length, value (TLV) structures. An LLDP frame should start with mandatory TLVs (e.g., Chassis ID, Port ID, and Time to live). These mandatory TLVs are followed by any number of optional TLVs. The frame should end with a special TLV named end of LLDPDU. The IEEE 802.1ab specification contains a description of all of the TLV types.

TKIP

TKIP is an acronym for **T**emporal **K**ey **I**ntegrity **P**rotocol. It used in WPA to replace WEP with a new encryption algorithm. TKIP comprises the same encryption engine and RC4 algorithm defined for WEP. The key used for encryption in TKIP is 128 bits and changes the key used for each packet.

Transparent clock

A device that measures the time taken for a Precision Time Protocol event message to transit the device and provides this information to clocks receiving this PTP event message.

Two-step clock

A clock that provides time information using the combination of an event message and a subsequent general message.

U**UDP**

UDP (**U**ser **D**atagram **P**rotocol) is a communications protocol that uses the Internet Protocol (**IP**) to exchange the messages between computers. UDP is an alternative to the Transmission Control Protocol (**TCP**) that uses the Internet Protocol (IP). Unlike TCP, UDP does not provide the service of dividing a message into packet datagrams, and UDP doesn't provide reassembling and sequencing of the packets. This means that the application program that uses UDP must be able to make sure that the entire message has arrived and is in the right order. Network applications that want to save processing time because they have very small data units to exchange may prefer UDP to TCP.

UDP provides two services not provided by the IP layer. It provides port numbers to help distinguish different user requests and, optionally, a checksum capability to verify that the data arrived intact.

Common network applications that use UDP include the Domain Name System (**DNS**), streaming media applications such as IPTV, Voice over IP (VoIP), and Trivial File Transfer Protocol (**TFTP**).

UNI

(User Network Interface) the physical interface or port that is the demarcation between the customer and the service provider/Cable Operator/Carrier/MSO per MEF 4.

UNI-C

UNI – Customer per MEF 4. A compound architectural component on the Subscriber side of the UNI that represents all the functions required to connect a subscriber to a MEN (per MEF 27).

UNI-N

UNI – Network per MEF 4. A compound functional element used to represent all of the functional elements required to connect a MEN to a MEN subscriber implementing a UNI C. The functional elements within the Customer Edge that supports the MEN Subscriber's technical capabilities and compliance to the UNI specification. A set of one or more functional elements that supports the MEN Service Provider's technical capabilities and compliance to the UNI specification. (Per MEF 27.)

Unicast address

In IPv6, an identifier for a single interface. A packet sent to a unicast address is delivered to the interface identified by that address.

UPnP

UPnP (**U**niversal **P**lug **a**nd **P**lay). The goals of 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.

Upper layer

In IPv6, a protocol layer immediately above IPv6. Examples are transport protocols such as TCP and UDP, control protocols such as ICMP, routing protocols such as OSPF, and internet or lower-layer protocols being "tunnelled" over (i.e., encapsulated in) IPv6 such as IPX, AppleTalk, or IPv6 itself.

User Priority

User Priority is a 3-bit field storing the priority level for the 802.1Q frame. It is also known as [PCP](#).

V**Valid IPv6 address**

In IPv6, a preferred or deprecated address. A valid address may appear as the source or destination address of a packet, and the internet routing system is expected to deliver packets sent to a valid address to their intended recipients. See the IETF "[Recommendation for IPv6 Address Text Representation](#)" or use a validator for IPv6 address formats such as <http://www.intermapper.com/ipv6validator>.

Valid lifetime

In IPv6, the length of time an address remains in the valid state (i.e., the time until invalidation). The valid lifetime must be greater than or equal to the preferred lifetime. When the valid lifetime expires, the address becomes invalid.

VeriPHY®

Cable diagnostics that detect cable conditions such as cable length, opens, shorts, coupling between pairs, and termination status. The VERIPHY trademark was assigned a serial number by the USPTO June 11, 2002 (Type Of Mark: Service Mark). The current federal status of this trademark filing is Cancelled - Section 8.

VLAN

Virtual LAN. A method to restrict communication between switch ports. **VLANs** can be used for the following applications:

VLAN unaware switching: This is the default configuration. All ports are VLAN unaware with Port **VLAN ID 1** and members of VLAN 1. This means that MAC addresses are learned in VLAN 1, and the switch does not remove or insert VLAN tags.

VLAN aware switching: This is based on the IEEE 802.1Q standard. All ports are VLAN aware. Ports connected to VLAN aware switches are members of multiple VLANs and transmit tagged frames. Other ports are members of one VLAN, set up with this Port VLAN ID, and transmit untagged frames.

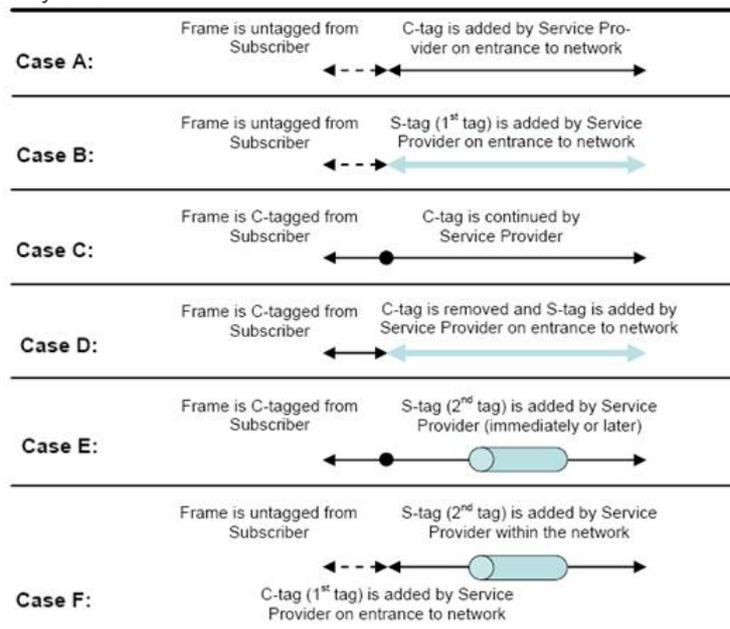
Provider switching: This is also known as Q-in-Q switching. Ports connected to subscribers are VLAN unaware, members of one VLAN, and set up with this unique Port VLAN ID. Ports connected to the service provider are VLAN aware, members of multiple VLANs, and set up to tag all frames. Untagged frames received on a subscriber port are forwarded to the provider port with a single VLAN tag. Tagged frames received on a subscriber port are forwarded to the provider port with a double VLAN tag.

VLAN ID

VLAN ID is a 12-bit field specifying the VLAN to which the frame belongs.

VLAN Tagging

In cases A - D below, a SOAM PDU is initiated by a Customer, and as it flows over the data path it continues to be processed and treated as a SOAM PDU. These frames exist in the OAM Flow Space seen by the Service Provider and the Operator. Thus MEG Levels used at any point can be seen by any other point in the path (subject to [IEEE 802.1ag restrictions on the extent of the MEG Levels). So different parties, such as the Service Provider and Operator, must coordinate the use of all levels that they share.



In Cases E and F above, the SOAM PDUs that were inserted in the un-tagged or single-tagged portions of the path are invisible to all points that are double tagged (since the double-tagged part of the path (the ‘tunnel’) has hidden the fact that a frame is a SOAM PDU with the addition of a second (outer) tag). These frames do not exist in the OAM Flow Space seen by the Service Provider and the Operator. Within the double-tagging, SOAM PDUs can be inserted and they can use any MEG Level without consideration for the MEG levels used by SOAM PDUs that use single tags.

Voice VLAN

Voice VLAN is VLAN configured specially for voice traffic. By adding the 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.

VoIP

Voice over Internet Protocol.

VTY

(Virtual Type Terminal) - A **vty** interface and password must be created in order to enable Telnet access to an IPv6 router. Also Virtual TTY (VTY).

W

WEP

WEP (**W**ired **E**quivalent **P**rivacy) is a deprecated algorithm to secure IEEE 802.11 wireless networks. Wireless networks broadcast messages using radio, so are more susceptible to eavesdropping than wired networks. When introduced in 1999, WEP was intended to provide confidentiality comparable to that of a traditional wired network (Wikipedia).

WiFi

WiFi (**W**ireless **F**idelity) is meant to be used generically when referring of any type of 802.11 network, whether 802.11b, 802.11a, dual-band, etc. The term is promulgated by the Wi-Fi Alliance.

WPA

WPA (**W**i-Fi **P**rotected **A**ccess) was created in response to several serious weaknesses researchers had found in the previous system , Wired Equivalent Privacy (WEP). WPA implements the majority of the IEEE 802.11i standard, and was intended as an intermediate measure to take the place of WEP while 802.11i was prepared. WPA is specifically designed to also work with pre-WPA wireless network interface cards (through firmware upgrades), but not necessarily with first generation wireless access points. WPA2 implements the full standard, but will not work with some older network cards (Wikipedia).

WPA-PSK

WPA-PSK (**W**i-Fi **P**rotected **A**ccess - **P**re **S**hared **K**ey) was designed to enhance the security of wireless networks. There are two flavours of WPA: enterprise and personal. Enterprise is meant for use with an IEEE 802.1X authentication server, which distributes different keys to each user. Personal WPA utilizes less scalable 'pre-shared key' (PSK) mode, where every allowed computer is given the same passphrase. In PSK mode, security depends on the strength and secrecy of the passphrase. The design of WPA is based on a Draft 3 of the IEEE 802.11i standard (Wikipedia)

WPA-Radius

WPA-Radius (**W**i-Fi **P**rotected **A**ccess - Radius) (802.1X authentication server) was designed to enhance the security of wireless networks. There are two flavours of WPA: enterprise and personal. Enterprise is meant for use with an IEEE 802.1X authentication server, which distributes different keys to each user. Personal WPA utilizes less scalable 'pre-shared key' (PSK) mode, where every allowed computer is given the same passphrase. In PSK mode, security depends on the strength and secrecy of the passphrase. The design of WPA is based on Draft 3 of the IEEE 802.11i standard (Wikipedia).

WTR

WTR (**W**ait **T**o **R**estore) is the time a fail on a resource has to be 'not active' before restoration back to this (previously failing) resource is done.

WTB

The **W**ait **T**o **B**lock (WTB) timer is employed by the RPL owner to delay reversion after a forced switch or manual switch has been cleared. From ITU-T Rec.G.8032/Y.1344 (03/2010).

SAT (Service Activation Testing) Terms

Bandwidth profile

In Ethernet SAT, provides measurements such as CIR, CBS, EIR, EBS, CF, and CM.

In carrier Ethernet, a characterization of ingress Service Frame arrival times and lengths at a reference point (UNI), and a specification of the disposition of each Service Frame - based on its level of compliance with the Bandwidth Profile. Per MEF 6.1, 10, 14, 15, 7, 9, 27, 29. Can be provided on the basis of:

Bandwidth profile per CoS ID: A bandwidth profile applied on a per-Class of Service basis.

Bandwidth profile per EVC: A bandwidth profile applied on a per-EVC basis.

Bandwidth profile per UNI: A bandwidth profile applied on a per-UNI basis.

Discovery

In Ethernet SAT, the procedure used by Ethernet Test Equipment to find ports or devices on the EVC which support the Latching Loopback feature.

ESCE

(Ethernet SAT Control Entry)

ESCF

(Ethernet Subscriber Conditioning Function) per MEF 12.1. The EtherSAT processing entity responsible for classification, filtering, metering, marking, policing, scheduling, shaping and, in general, conditioning the subscriber flow into and out of a UNI-N. See MEF 12, 17.

ETA

Ethernet Test Application per MEF SAT. A software module resident in a Network Element that enables the Network Element to perform Service Activation Testing and activate/deactivate loopback devices (per MEF SAT).

ETH

Ethernet Test Head per MEF SAT. A fixed, embedded Ethernet device which may include an Ethernet Generator, Ethernet Collector, and/or Latching Loopback function which is dedicated for Ethernet testing (per MEF SAT).

Ethernet Equipment

A network element or test equipment installed or used in an Ethernet Network per “Technical Specification MEF x.0, Latching Loopback Protocol and Functionality, Version 011”.

Ethernet Test Equipment

A general term to include a Ethernet Test Application, Ethernet Test Head and/or ETI (per MEF SAT).

ETI

(Ethernet Test Instrument) - A portable, external Ethernet testing device not permanently installed in the network which may include an Ethernet Generator, Ethernet Collector, and/or Loopback function(s) (per MEF SAT).

External Facing Loopback

A loopback in a device where the frames pass through the minimum possible functional blocks of the device before they are rerouted to the other direction where they again pass through the minimum possible functional blocks of the device. Per “Technical Specification MEF x.0, Latching Loopback Protocol and Functionality, Version 011”. Contrast “Internal Facing Loopback”.

EtherSAM

See “Y.1564”.

FDV

(Frame Delay Variation) - one of several SAT Performance criteria. Per MEF 10.2, 14, 17, 19. (Often misquoted as ‘Jitter’.) The difference in delay of two Service Frames, per MEF 10.1. Replaced by Frame Delay Variation Performance in MEF 10.2.

FLR

Frame Loss Ratio - one of several SAT Performance criteria. Per MEF 10.1, 14, 17, 19.

Frame Delay Variation Performance

A measure of the variation in the delays experienced by different Service Frames belonging to the same CoS instance. See also MEF 10.2, 14, 15, and 19.

Frame Loss Ratio Performance

Frame Loss Ratio is a measure of the number of lost frames between the ingress UNI and the egress UNI. Frame Loss Ratio is expressed as a percentage. Per MEF 10.2, 19, 10, 14, 15.

FRE

(Frame Reference Event)

FTD

(Frame Transfer Delay) - One of several EtherSAT Performance criteria.

Information rate

The average bit rate of Ethernet service frames at the measurement point starting with the first MAC address bit and ending with the last FCS bit. See MEF 10.2 for a clear discussion of Ethernet service frames and specific examples of information rates (the CIR and EIR). For example, a 100 Mbps Ethernet port can handle a total information rate of about 77 Mbps to 99 Mbps, depending on the average frame size of the transmitted Ethernet frames.

Internal Facing Loopback

In Ethernet SAT, a loopback in a device where the frames pass through all the possible functional blocks of the device before they are rerouted to the other direction where they again pass through all the possible functional blocks of the device. Contrast “External Facing Loopback”.

LL (Latching Loopback)

In Ethernet SAT, a configured function within an Ethernet Equipment where frames are returned to the entity which sent them. Per “Technical Specification MEF x.0, Latching Loopback Protocol and Functionality, Version 011”.

LLD (Latching Loopback Device)

In Ethernet SAT, any NE or Ethernet Test Equipment that supports the Latching Loopback function and protocol defined in “Technical Specification MEF x.0, Latching Loopback Protocol and Functionality, Version 011”.

Loopable Frame

In Ethernet SAT, a frame eligible for loopback because it **a)** belongs to the same EVC that is currently in the Latching Loopback Active state, and **b)** is not required to be discarded for any particular reason listed in the MEF numbered requirements. Per “Technical Specification MEF x.0, Latching Loopback Protocol and Functionality, Version 011”.

MP

(Measurement Point) - In Ethernet SAT, the boundary between a bridge and an adjacent exchange link where performance reference events can be observed and measured. A section or a combination of sections is measurable if bounded by a set of MPs. See the Rec. ITU-T Y.1564 (03/2011) for the sections that are measurable.

Non-Test Side

In Ethernet SAT, the Latching Loopback Port which faces away from the test equipment that invokes the Latching Loopback on the EVC. Few EVC frames egress the Non-Test Side. Non-Test Side EVC Ingress frames do not egress the Test Side. Per “Technical Specification MEF x.0, Latching Loopback Protocol and Functionality, Version 011”. Contrast “Test Side” below:

Performance criteria

In Ethernet SAT, measurements such as FTD, FDV, FLR, AVAIL, etc.

SA/EVC

In Ethernet SAT, frames that have the same MAC Source Address and are a part of the same Ethernet Virtual Connection. Per “Technical Specification MEF x.0, Latching Loopback Protocol and Functionality, Version 011”.

SAC

(Service Acceptance Criteria) - In Ethernet SAT, a set of criteria used to ensure that a service meets its functional and quality requirements, and that it is ready to operate when deployed.

SAT

(Service Activation Testing) per MEF SAT.

Service activation

In Ethernet SAT, the step of bringing a network feature into operation or eventual customer use, before that customer is notified that the feature is ready to use.

Test Side

In Ethernet SAT, the side of the Latching Loopback Device which includes the port used to communicate with the test equipment regarding Latching Loopback functions. Per “Technical Specification MEF x.0, Latching Loopback Protocol and Functionality, Version 011”. Contrast “Non-Test Side” above.

ULR

Utilized Line Rate, used in Ethernet SAT.

WRED

(Weighted Random Early Detection) generally drops packets selectively based on IP precedence. Packets with a higher IP precedence are less likely to be dropped than packets with a lower precedence. Thus higher priority traffic is delivered with a higher probability than lower priority traffic. allows configuring different drop profiles to different traffic flows hence providing different QoS for different types of traffic. WRED is able to distinguish traffic flows by examining IP precedence value (TOS field in IP header) or in case of differentiated service enabled flow, the DSCP.

Y.1564

ITU-T Y.1564 (EtherSAM) is an Ethernet service activation test methodology. The Y.1564 recommendation defines a test methodology for use in assessing the proper configuration and performance of an Ethernet network to deliver Ethernet-based services. The Y.1564 out-of-service test methodology was created so service providers have a standard way of measuring the performance of Ethernet-based services. ITU-T Y.1564 is the new ITU-T standard for turning up, installing, and troubleshooting Ethernet-based services. It is the only standard test methodology that allows for complete validation of Ethernet service-level agreement's (SLAs) in a single test. ITU-T Y.1564 is designed around three key objectives:

1. To serve as a network service level agreement (SLA) validation tool, ensuring that a service meets its guaranteed performance settings in a controlled test time.
2. To ensure that all services carried by the network meet their SLA objectives at their maximum committed rate, proving that under maximum load network devices and paths can support all the traffic as designed.
3. To perform medium- and long-term service testing, confirming that network element can properly carry all services while under stress during a soaking period.

ITU-T Y.1564 defines an out-of-service test methodology to assess the proper configuration and performance of an Ethernet service prior to customer notification and delivery. The test methodology applies to point-to-point and point-to-multipoint connectivity in the Ethernet layer and to the network portions that provide, or contribute to, the provisioning of such services. This recommendation does not define Ethernet network architectures or services, but rather defines a methodology to test Ethernet-based services at the service activation stage. In particular, it is aimed at improving on RFC 2544 functionality.

10 Gigabit Ethernet Terms

10 gigabit Ethernet

(10GE or 10GbE or 10 GigE) refers to various technologies for transmitting Ethernet frames at a rate of 10 gigabits per second (10 billion bits per second), first defined by the IEEE 802.3ae-2002 standard. Like previous versions of Ethernet, 10GbE supports both copper and fibre cabling. However, due to its higher bandwidth requirements, higher-grade copper cables are required: category 6a or Class F/Category 7 cables for links up to 100m. Unlike previous Ethernet standards, 10 gigabit Ethernet only defines full duplex point-to-point links which are generally connected by network switches. Half duplex operation does not exist in 10GbE.

The 10 gigabit Ethernet standard encompasses a number of different physical layer (PHY) standards. A networking device may support different PHY types through pluggable PHY modules, such as those based on SFP+. Over time, market forces will determine the most popular 10GE PHY types. At the time that the 10 gigabit Ethernet standard was developed, interest in 10GbE as a wide area network (WAN) transport led to the introduction of a WAN PHY for 10GbE. The WAN PHY encapsulates Ethernet packets in SONET OC-192c frames and operates at a slightly slower data-rate (9.95328 Gbps) than the local area network (LAN) PHY. Both share the same physical medium-dependent sublayers so can use the same optics. The WAN PHY can support maximum link distances up to 80 km depending on the fibre standard employed.

Backplane Ethernet

Backplane - also known by its task force name 802.3ap - is used in backplane applications such as blade servers and routers/switches with upgradable line cards. 802.3ap implementations are required to operate in an environment comprising up to 1 meter (39 in) of copper printed circuit board with two connectors. The standard defines two port types for 10 Gbit/s (10GBASE-KX4 and 10GBASE-KR) and a 1 Gbit/s port type (1000BASE-KX). It also defines an optional layer for FEC, a backplane autonegotiation protocol and link training for 10GBASE-KR where the receiver can set a three tap transmit equalizer. The autonegotiation protocol selects between 1000BASE-KX, 10GBASE-KX4, 10GBASE-KR or 40GBASE-KR4 operation. 40GBASE-KR4 is defined in 802.3ba. New backplane designs use 10GBASE-KR rather than 10GBASE-KX4.

10GBASE-KX4

Operates over four backplane lanes and uses the same physical layer coding (defined in IEEE 802.3 Clause 48) as 10GBASE-CX4.

10GBASE-KR

Operates over a single backplane lane and uses the same physical layer coding (defined in IEEE 802.3 Clause 49) as 10GBASE-LR/ER/SR.

10GBASE-T

10GBASE-T, or IEEE 802.3an-2006, is a standard released in 2006 to provide 10 Gbit/s connections over unshielded or shielded twisted pair cables, over distances up to 100 meters (330 ft). Although category 6a is required to reach the full 100 meters (330 ft), category 5e is good for up to 45 meters (148 ft) and category 6 will reach 55 meters (180 ft).[22] 10GBASE-T cable infrastructure can also be used for 1000BASE-T allowing a gradual upgrade from 1000BASE-T using autonegotiation to select which speed to use. 10GBASE-T has latency in the range 2 to 4 microseconds compared to 1 to 12 microseconds on 1000BASE-T. As of 2010 10GBASE-T silicon is available from several manufacturers with claimed power dissipation of 3-4 W at structure widths of 40 nm. With 28 nm in development, power will continue to decline.

10GBASE-T uses the IEC 60603-7 8P8C (commonly known as RJ45) connectors already widely used with Ethernet. Transmission characteristics are now specified to 500 MHz. To reach this frequency Category 6A or better balanced twisted pair cables specified in ISO/IEC 11801 amendment 2 or ANSI/TIA-568-C.2 are needed to carry 10GBASE-T up to distances of 100 m. Category 6 cables can

carry 10GBASE-T for shorter distances when qualified according to the guidelines in ISO TR 24750 or TIA-155-A.

10GBASE-SR

10GBASE-SR ("short range") is a port type for multi-mode fibre and uses 850 nm lasers. Its Physical Coding Sublayer 64b/66b PCS is defined in IEEE 802.3 Clause 49 and its Physical Medium Dependent PMD in Clause 52. It delivers serialized data at a line rate of 10.3125 Gbit/s.

Over obsolete FDDI-grade 62.5 micrometres multimode fibre cabling it has a maximum range of 26 meters. Over 62.5 micrometres OM1 it has a range of 33 meters; over 50 micrometres OM2 a range of 82 meters; over OM3 300 meters and over OM4 400 meters. OM3 and OM4 are the preferred choices for structured optical cabling within buildings. MMF has the advantage over SMF of having lower cost connectors because of its wider core.

There is a non-standard lower cost, lower power variant sometimes referred to as 10GBASE-SRL (10GBASE-SR lite). This is inter-operable with 10GBASE-SR but only has a reach of 100 meters.

10GBASE-LR

10GBASE-LR ("long reach") is a port type for single-mode fibre and uses 1310 nm lasers. Its Physical Coding Sublayer 64b/66b PCS is defined in IEEE 802.3 Clause 49 and its Physical Medium Dependent PMD in Clause 52. It delivers serialized data at a line rate of 10.3125 Gbit/s.

10GBASE-LR has a specified reach of 10 kilometres (6.2 mi), but 10GBASE-LR optical modules can often manage distances of up to 25 kilometres (16 mi) with no data loss.

10GBASE-LRM

10GBASE-LRM, (Long Reach Multimode) originally specified in IEEE 802.3aq is a port type for multimode fibre and uses 1310 nm lasers. Its Physical Coding Sublayer 64b/66b PCS is defined in IEEE 802.3 Clause 49 and its Physical Medium Dependent PMD in Clause 68. It delivers serialized data at a line rate of 10.3125 Gbit/s.

10GBASE-LRM supports distances up to 220 meters (720 ft) on FDDI-grade multimode fibre and the same 220m maximum reach on OM1, OM2 and OM3 fibre types. 10GBASE-LRM reach is not quite as far as the older 10GBASE-LX4 standard.

To ensure that specifications are met over FDDI-grade, OM1 and OM2 fibres, the transmitter should be coupled through a mode conditioning patch cord. No mode conditioning patch cord is required for applications over OM3 or OM4.

Some 10GBASE-LRM transceivers also support distances up to 300 meters (980 ft) on standard single-mode fibre (SMF, G.652), however this is not part of the IEEE or MSA specification. 10GBASE-LRM uses electronic dispersion compensation (EDC) for receive equalization.

10GBASE-LRM has been a failure in the market.

10GBASE-ER

10GBASE-ER ("extended reach") is a port type for single-mode fibre and uses 1550 nm lasers. Its Physical Coding Sublayer 64b/66b PCS is defined in IEEE 802.3 Clause 49 and its Physical Medium Dependent PMD in Clause 52. It delivers serialized data at a line rate of 10.3125 Gbit/s.

The 10GBASE-ER transmitter is implemented with an externally modulated laser (EML).

10GBASE-ER has a reach of 40 kilometres (25 mi) over engineered links and 30 km over standard links.

10GBASE-ZR

Several manufacturers have introduced 80 km (50 mi) range ER pluggable interfaces under the name 10GBASE-ZR. This 80 km PHY is not specified within the IEEE 802.3ae standard and manufacturers have created their own specifications based upon the 80 km PHY described in the OC-192/STM-64 SDH/SONET specifications. The 802.3 standard will not be amended to cover the ZR PHY.

802.3 Standards for 10GbE

Over the years the IEEE 802.3 working group has published several standards relating to 10GbE. These included: 802.3ae-2002 (fibre -SR, -LR, -ER and -LX4 PMDs), 802.3ak-2004 (-CX4 copper twin-ax InfiniBand type cable), 802.3an-2006 (10GBASE-T copper twisted pair), 802.3ap-2007 (copper backplane -KR and -KX4 PMDs) and 802.3aq-2006 (fibre -LRM PMD with enhanced equalization). The

802.3ae-2002 and 802.3ak-2004 amendments were consolidated into the IEEE 802.3-2005 standard. IEEE 802.3-2005 and the other amendments were consolidated into IEEE Std 802.3-2008.

ALR

(Automatic Link Restoration) After a link failure condition has been corrected, the device will automatically re-establish the link in all network conditions using the ALR feature.

Cat 6 (Category 6) Cable

Category 6 cable, commonly referred to as Cat 6, is a standardized cable for Gigabit Ethernet and other network physical layers that is backward compatible with the Category 5/5e and Category 3 cable standards. Compared to Cat 5 and Cat 5e, Cat 6 provides more stringent specifications for crosstalk and system noise. The Cat 6 cable standard provides performance of up to 250 MHz and is suitable for 10BASE-T, 100BASE-TX (Fast Ethernet), 1000BASE-T/1000BASE-TX (Gigabit Ethernet) and 10GBASE-T (10-Gigabit Ethernet). Category 6 cable has a reduced maximum length when used for 10GBASE-T. Like most earlier twisted-pair cable, Category 6 cable contains four twisted wire pairs. Attenuation, near end crosstalk (NEXT), and PSNEXT (power sum NEXT) in Cat 6 cable and connectors are all much lower than Cat 5 or Cat 5e, which uses 24 AWG wire. The increase in performance with Cat 6 comes mainly from increased (22 AWG) wire size. Because the conductor sizes are generally the same, Cat 6 jacks may also be used with Cat 5e cable.

Category 6 cable can be identified by the printing on the side of the cable sheath. Cat 6 patch cables are normally terminated in 8P8C modular connectors. If Cat 6 rated patch cables, jacks, and connectors are not used with Cat 6 wiring, overall performance is degraded to that of the cable or connector.

Connectors use either T568A or T568B pin assignments; although performance is comparable provided both ends of a cable are the same, T568B is a deprecated standard in the US and no longer supported by TIA.

Category 6a Cable (Augmented Category 6)

Category 6a cable, or Augmented Category 6, is characterized to 500 MHz and has improved alien crosstalk characteristics, allowing 10GBASE-T to be run for the same distance as previous protocols. The latest standard from the TIA for enhanced performance standards for twisted pair cable systems was defined in February 2008 in ANSI/TIA/EIA-568-B.2-10. Category 6a is defined at frequencies up to 500 MHz—twice that of Cat. 6.

Category 6a performs at improved specifications, in particular in the area of alien crosstalk, as compared to Cat 6 UTP (unshielded twisted pair), which exhibited high alien noise in high frequencies. The global cabling standard ISO/IEC 11801 has been extended by the addition of amendment 2, which defines new specifications for Cat 6A components and Class EA permanent links. These new global Cat 6A/Class EA specifications require a new generation of connecting hardware, which offer superior performance compared to existing products based on the American TIA standard.

Note the performance difference between ISO/IEC and EIA/TIA component specifications for the NEXT transmission parameter. At a frequency of 500 MHz, an ISO/IEC Cat 6A connector performs 3 dB better than a Cat 6A connector that conforms with the EIA/TIA specification. The 3 dB represents a 100% increase of near-end crosstalk noise reduction when measured in absolute magnitudes.

When used for 10GBASE-T, Cat 6 cable's maximum length is 55 meters (180 ft) in a favourable alien crosstalk environment, but only 37 meters (121 ft) in a hostile alien crosstalk environment, such as when many cables are bundled together. However, because the effects of alien crosstalk environments on cables are difficult to determine prior to installation, it is highly recommended that all Cat 6 cables being used for 10GBASE-T are electrically tested once installed. With its improved specifications, Cat6 A does not have this limitation and can run 10GBASE-T at 100 meters (330 ft) without electronic testing.

LRM

Long Reach Multimode. See "10GBASE-LRM".

MSA

Multi-Source Agreements. enhanced Small Form-factor Pluggable transceiver. To support different 10GbE physical layer standards, many interfaces consist of a standard socket into which different PHY modules may be plugged. Physical layer modules are not specified in an official standards body but by

multi-source agreements (MSAs) that can be negotiated more quickly. Relevant MSAs for 10GbE include XENPAK (and related X2 and XPAK), XFP and SFP+. When choosing a PHY module, a designer considers cost, reach, media type, power consumption, and size (form factor). A single point-to-point link can have different MSA pluggable formats on either end (e.g. XPAK and SFP+) as long as the 10GbE optical or copper interface (e.g. 10GBASE-SR) inside the pluggable is identical. See also “SFP+”.

SFP+

enhanced Small Form-factor Pluggable transceiver. To support different 10GbE physical layer standards, many interfaces consist of a standard socket into which different PHY modules may be plugged. Physical layer modules are not specified in an official standards body but by multi-source agreements (MSAs) that can be negotiated more quickly. Relevant MSAs for 10GbE include XENPAK (and related X2 and XPAK), XFP and SFP+. When choosing a PHY module, a designer considers cost, reach, media type, power consumption, and size (form factor). A single point-to-point link can have different MSA pluggable formats on either end (e.g. XPAK and SFP+) as long as the 10GbE optical or copper interface (e.g. 10GBASE-SR) inside the pluggable is identical. See also “MSA”.

The newest module standard is the enhanced small form-factor pluggable transceiver, generally called SFP+. Based on the small form-factor pluggable transceiver (SFP) and developed by the ANSI T11 fibre channel group, it is smaller still and lower power than XFP. SFP+ has become the most popular socket on 10GE systems. SFP+ modules do only optical to electrical conversion, no clock and data recovery, putting a higher burden on the host's channel equalization. SFP+ modules share a common physical form factor with legacy SFP modules, allowing higher port density than XFP and the re-use of existing designs for 24 or 48 ports in a 19" rack width blade.

SFP+ modules can further be grouped into two types of host interfaces: linear or limiting. Limiting modules are preferred except when using old fibre infrastructure which requires the use of the linear interface provided by 10GBASE-LRM modules.

TLPT

(Transparent Link Pass Through) will notify an end device of a link failure just like Link Pass Through; however, it uses a different method for “passing through” this information. Transparent Link Pass Through sends a link loss signal over the fibre, instructing the remote converter to shut down the copper port thus notifying the end device, while maintaining the fibre link between the two converters. With TLPT, an End device automatically notified of link loss, and the Fibre link remains up as it carries a link loss signal

XAUI

XAUI is a standard for extending the XGMII (10 Gigabit Media Independent Interface) between the MAC and PHY layer of 10 Gigabit Ethernet (10GbE). XAUI is pronounced "zowie", a concatenation of the Roman numeral X, meaning ten, and the initials of "Attachment Unit Interface". The XGMII Extender, which is composed of an XGXS at the MAC end, an XGXS at the PHY end and a XAUI between them, is to extend the operational distance of the XGMII and to reduce the number of interface signals.

Applications include extending the physical separation possible between MAC and PHY components in a 10 Gigabit Ethernet system distributed across a circuit board.

Sync-E Terms

ACR

The Adaptive Clock Recovery technique does not require the support of a network-wide synchronization signal to regenerate the timing. In this case, the timing recovery process is based on the (inter-)arrival time of the packets (e.g., timestamps or CES packets). The information carried by the packets could be used to support this operation. Two-way or one-way protocols can be used. NOTE – For these purposes, this definition relates to frequency synchronization only.

PSN

Packet –Switched Networks

PDH

Plesiochronous Digital Hierarchy

SDH

Synchronous Digital Hierarchy

SONET

Synchronous Optical NETWORK

CES

Circuit Emulation Services

CES IWF

The set of functions within the IWF that supports the service clock domain. This includes the function to recover the service clock timing.

CES island

Segment of a network, based on packet-switched technologies, that emulates either the characteristics of a circuit-switched network or of a PDH/SDH transport network, in order to carry CBR services (e.g., E1).

Free running mode

An operating condition of a clock, the output signal of which is strongly influenced by the oscillating element and not controlled by servo phase-locking techniques. In this mode the clock has never had a network reference input, or the clock has lost external reference and has no access to stored data, that could be acquired from a previously connected external reference. Free-run begins when the clock output no longer reflects the influence of a connected external reference, or transition from it. Free-run terminates when the clock output has achieved lock to an external reference. Contrast “Holdover mode” and “Locked mode”.

Holdover mode

An operating condition of a clock which has lost its controlling reference input and is using stored data, acquired while in locked operation, to control its output. The stored data are used to control phase and frequency variations, allowing the locked condition to be reproduced within specifications. Holdover begins when the clock output no longer reflects the influence of a connected external reference, or transition from it. Holdover terminates when the output of the clock reverts to locked mode condition. Contrast “Free running mode” and “Locked mode”.

Frequency source traceability

A relationship where the frequency of all clocks in a system is referenced back to a single physical clock. Under normal operating conditions, all clocks will have the same average frequency in a source-traceable system. Thus, the phase error or maximum time interval error (MTIE) between all clocks in such a system is bounded. Note: a different case exists when the clocks have frequency traceability towards accurate master clocks (not necessarily the same pieces of equipment). This is connected to the concept of plesiochronous as defined in b-ITU-T G.810. An example is when the clocks have the "PRC traceability" (i.e., traceability to ITU-T G.811 clocks) in a synchronization network based on the distributed PRC architecture.

Locked mode

An operating condition of a slave clock in which the output signal is controlled by an external input reference such that the clock's output signal has the same long-term average frequency as the input reference, and the time error function between output and input is bounded. Locked mode is the expected mode of operation of a slave clock. Contrast "Free running mode" and "Holdover mode".

Network clock

The clock generating the network clock signal.

Network clock domain

The set of functions dedicated to support the synchronization network (network clock).

Network clock signal

A reference timing signal that is used as a reference to allow mapping and demapping of a service clock at ingress and egress points of the network respectively. In some applications, the signal could be asynchronous and generated by free running clocks with low requirements in terms of frequency accuracy (e.g., in the Ethernet network, the physical layer can operate up to ± 100 ppm). In other applications, an accurate reference timing signal is needed. In this case, the signal is typically traceable to a PRC under normal conditions, and the distribution of this signal across the network is accomplished by a synchronization network. Note: for these purposes, it is assumed that a properly high accurate signal is always involved.

Due to that, the network clock signal definition can be considered to coincide with the definition of synchronization network clock signal, and the two terms are used interchangeably.

Network-synchronous operation

Synchronization of the physical layer (usually by a timing distribution of a timing signal traceable to a primary reference clock (PRC), see ITU-T G.811).

PNT-F

Packet network timing function; the set of functions within the IWF that supports the synchronization network clock domain (see Figure B.2). This includes the function to recover and distribute the timing carried by the synchronization network. The PNT-F clocks may be part of the IWF or may be part of any other network element in the packet network. When the PNT-Fs are part of the IWF, they may support the CES IWF and/or change the layer over which timing is carried (i.e., from packet to physical layer and vice versa).

Service clock

The clock generating the service clock signal.

Service clock domain

The set of functions dedicated to support the CES timing function (service clock).

Service clock signal

The timing information that is associated with a specific service supported by a network. For instance, in case of E1 TDM service, the timing shall be 2048 kbit/s ± 50 ppm.

SMB

(SubMiniature version B) coaxial RF connectors developed in the 1960s. SMB connectors are smaller than SMA connectors. SMB connectors feature a snap-on coupling and are available in either 50 Ω or 75 Ω impedance. They offer excellent electrical performance from DC to 4 GHz. An SMB jack has a male centre pin, while an SMB plug has a female basket. Connectors are available for two SMB cable sizes: 2.6/50+75 S (3 mm outer / 1.7 mm inner diameter) and 2/50 S (2.2 mm outer / 1 mm inner diameter).

Synchronization network clock

The equipment that provides the timing signal in the synchronization network.

Synchronization network clock signal

The reference timing signal distributed by the synchronization network. This signal is traceable to an accurate master (e.g., PRC).

TDM

Time division multiplex; conventionally refers to the isochronous bit streams used in telephony networks; in particular, those belonging to PDH (plesiochronous digital hierarchy) as described in ITU-T G.705. The bit rates traditionally used in various regions of the world are detailed in ITU-T G.702. Examples of the signals covered by the TDM definition are those belonging to PDH and SDH hierarchies.

Stabilization period

The period beginning at the point in time when a validated timing source has been selected by the IWF and ending when the output timing characteristics are within the output jitter and wander requirements. Reference Rec. ITU-T G.8261/Y.1361 (04/2008).

Wander budget

Wander generated at the output of a network island when an ideal reference timing signal is the input at the first network element of this network island.

Index

AAA Configuration	134	
Access Management Configuration		59
ACE Configuration	101, 109	
ACL Configuration	109	
ACL Port Configuration	101	
ACL Rate Configuration	104	
Aggregation	138	
Aggregation Configuration		137
APS	192	
ARP Inspection Configuration		130
Authentication Configuration		134
Authentication Configuration		53
Authentication password		73
Auto Configuration		30
change password		47, 48
change privilege level		50
Configuration Menu		23
configure Privilege Level		46, 48
configure User Name		46
DHCP Configuration		122
DPL		351
DSA		57
Dying Gasp		411
Engine ID		69
EPS Configuration		189, 191
EPS Instance Command		194
EPS Instance Configuration		192
EPS Instance State		195
EPS Protection Type		192
ERPS Configuration		225
Eth Services Configuration		272
EVC Configuration		279
Fast Leave		175
HTTPS Certificate Load		58
HTTPS Configuration		56
Immediate Leave		168
IP Configuration		25
IPv6 Configuration		29
LACP Configuration		141
Link Aggregation		138
Link OAM Configuration		144
LLDP Configuration		186
LOAM Configuration		144
Loop Protection Configuration		150
Loopback Configuration		207
MAC Address Table Configuration		244
MAC Table Learning		245
MEP Configuration		197
MIBs supported		643
MIP		200
Mirroring Configuration		353
MLD Snooping		180
Monitor Configuration		389
Monitoring Performance		214
MSTP		154
MVR Configuration		165
NAS Configuration		91
NTP Configuration		34
Port Aggregation		138
Port Configuration		42
Port Isolation		260
Private VLAN Configuration		258
Privilege Level Configuration		50
Provider Bridging		257
PVLAN Membership		258
QoS Configuration		309
Read Community		68
RED		351
Returns, product		636
RMON Configuration		79
RPL Configuration		235
RSA		57
RSTP		154
SNMP Configuration		62
SNMP Security Model		75, 77
SNMP Trap Configuration		68
SNMP v1 Traps		67
SNMP v2 Traps		67
SNMP v3 Traps		67
Snooping Configuration		122
SOAM Configuration		197
Spanning Tree Configuration		154
SSH Configuration		54
STP		154
STP Bridge Configuration		155
Sub-Ring Configuration		236
SysLog Configuration		41
Throttling		175
Trap Authentication		69
Trap Configuration		69
User Configuration		46
User Privilege Level		46
VLAN Configuration		252
VLAN ID		26
VLAN Translation Configuration		248
Write Community		69



Net2Edge Limited
Kulite House,
Stroudley Road,
Basingstoke
RG24 8UG
United Kingdom
Tel: +44 345 0130030