39XX/51XX Service Delivery, Aggregation and Virtualization Switches

Base Configuration

SAOS 6.18

What's inside...

009-3297-008 - Standard Revision A
January 2019

LEGAL NOTICES

Security

Ciena® cannot be responsible for unauthorized use of equipment and will not make allowance or credit for unauthorized use or access.

Contacting Ciena

Corporate Headquarters	410-694-5700 or 800-921-1144	www.ciena.com
Customer Technical Support/Warranty		
In North America	1-800-CIENA-24 (243-6224) 410-865-4961	
In Europe, Middle East, and Africa	800-CIENA-24-7 (800-2436-2247) +44-207-012-5508 00 0800 77 454 (Slovenia)	
In Asia-Pacific	800-CIENA-24-7 (800-2436-2247) +81-3-6367-3989 +91-124-4340-600 120 11104 (Vietnam) 000 8004401369 (India)	
In Caribbean and Latin America	800-CIENA-24-7 (800-2436-2247) 1230-020-0845 (Chile) 009 800-2436-2247 (Colombia) 0800-77-454 (Mexico and Peru) 00 008000442510 (Panama)	
Sales and General Information	North America: 1-800-207-3714 International: +44 20 7012 5555	E-mail: sales@ciena.com
In North America	410-694-5700 or 800-207-3714	E-mail: sales@ciena.com
In Europe	+44-207-012-5500 (UK)	E-mail: sales@ciena.com
In Asia	+81-3-3248-4680 (Japan)	E-mail: sales@ciena.com
In India	+91-22-42419600	E-mail: sales@ciena.com
In Latin America	011-5255-1719-0220 (Mexico City)	E-mail: sales@ciena.com
Training		E-mail: learning@ciena.com

For additional office locations and phone numbers, please visit the Ciena web site at www.ciena.com.

READ THIS LICENSE AGREEMENT (“LICENSE”) CAREFULLY BEFORE INSTALLING OR USING CIENA SOFTWARE OR DOCUMENTATION. THIS LICENSE IS AN AGREEMENT BETWEEN YOU AND CIENA COMMUNICATIONS, INC. (OR, AS APPLICABLE, SUCH OTHER CIENA CORPORATION AFFILIATE LICENSOR) (“CIENA”) GOVERNING YOUR RIGHTS TO USE THE SOFTWARE. BY INSTALLING OR USING THE SOFTWARE, YOU ACKNOWLEDGE THAT YOU HAVE READ THIS LICENSE AND AGREE TO BE BOUND BY IT.

1. License Grant. Ciena may provide **“**Software**”** to you either (1) embedded within or running on a hardware product or (2) as a standalone application, and Software includes upgrades acquired by you from Ciena or a Ciena authorized reseller. Subject to these terms, and payment of all applicable License fees including any usage-based fees, Ciena grants you, as end user, a non-exclusive, non-transferable, personal License to use the Software only in object code form and only for its intended use as evidenced by the applicable product documentation. Unless the context does not permit, Software also includes associated documentation.

2. Open Source and Third Party Licenses. Software excludes any open source or third-party programs supplied by Ciena under a separate license, and you agree to be bound by the terms of any such license. If a separate license is not provided, any open source and third party programs are considered “Software” and their use governed by the terms of this License.

3. Title. You are granted no title or ownership rights in or to the Software. Unless specifically authorized by Ciena in writing, you are not authorized to create any derivative works based upon the Software. Title to the Software, including any copies or derivative works based thereon, and to all copyrights, patents, trade secrets and other intellectual property rights in or to the Software, are and shall remain the property of Ciena and/or its licensors. Ciena's licensors are third party beneficiaries of this License. Ciena reserves to itself and its licensors all rights in the Software not expressly granted to you.

4. Confidentiality. The Software contains trade secrets of Ciena. Such trade secrets include, without limitation, the design, structure and logic of individual Software programs, their interactions with other portions of the Software, internal and external interfaces, and the programming techniques employed. The Software and related technical and commercial information, and other information received in connection with the purchase and use of the Software that a reasonable person would recognize as being confidential, are all confidential information of Ciena (“Confidential Information”).

5. Obligations. You shall:

i) Hold the Software and Confidential Information in strict confidence for the benefit of Ciena using your best efforts to protect the Software and Confidential Information from unauthorized disclosure or use, and treat the Software and Confidential Information with the same degree of care as you do your own similar information, but no less than reasonable care;

ii) Keep a current record of the location of each copy of the Software you make;

iii) Use the Software only in accordance with the authorized usage level;

iv) Preserve intact any copyright, trademark, logo, legend or other notice of ownership on any original or copies of the Software, and affix to each copy of the Software you make, in the same form and location, a reproduction of the copyright notices, trademarks, and all other proprietary legends and/or logos appearing on the original copy of the Software delivered to you; and

v) Issue instructions to your authorized personnel to whom Software is disclosed, advising them of the confidential nature of the Software and provide them with a summary of the requirements of this License.

6. Restrictions. You shall not:

i) Use the Software or Confidential Information a) for any purpose other than your own internal business purposes; and b) other than as expressly permitted by this License;

ii) Allow anyone other than your authorized personnel who need to use the Software in connection with your rights or obligations under this License to have access to the Software;

iii) Make any copies of the Software except such limited number of copies, in machine readable form only, as may be reasonably necessary for execution in accordance with the authorized usage level or for archival purposes only;

iv) Make any modifications, enhancements, adaptations, derivative works, or translations to or of the Software;

v) Reverse engineer, disassemble, reverse translate, decompile, or in any other manner decode the Software;

vi) Make full or partial copies of the associated documentation or other printed or machine-readable matter provided with the Software unless it was supplied by Ciena in a form intended for reproduction;

vii) Export or re-export the Software from the country in which it was received from Ciena or its authorized reseller unless authorized by Ciena in writing; or

viii) Publish the results of any benchmark tests run on the Software.

7. Audit: Upon Ciena's reasonable request you shall permit Ciena to audit the use of the Software to ensure compliance with this License.

8. U.S. Government Use. The Software is provided to the Government only with restricted rights and limited rights. Use, duplication, or disclosure by the Government is subject to restrictions set forth in FAR Sections 52-227-14 and 52-227-19 or DFARS Section 52.227-7013(C)(1)(ii), as applicable. The Software and any accompanying technical data (collectively “Materials”) are commercial within the meaning of applicable Federal acquisition regulations. The Materials were developed fully at private expense. U.S. Government use of the Materials is restricted by this License, and all other U.S. Government use is prohibited. In accordance with FAR 12.212 and DFAR Supplement 227.7202, the Software is commercial computer software and the use of the Software is further restricted by this License.

9. Term of License. This License is effective until the applicable subscription period expires or the License is terminated. You may terminate this License by giving written notice to Ciena. This License will terminate immediately if (i) you breach any term or condition of this License or (ii) you become insolvent, cease to carry on business in the ordinary course, have a receiver appointed, enter into liquidation or bankruptcy, or any analogous process in your home country. Termination shall be without prejudice to any other rights or remedies Ciena may have. Upon any termination of this License you shall destroy and erase all copies of the Software in your possession or control, and forward written certification to Ciena that all such copies of Software have been destroyed or erased. Your obligations to hold the Confidential Information in confidence, as provided in this License, shall survive the termination of this License.

10. Compliance with laws. You agree to comply with all laws related to your installation and use of the Software. Software is subject to U.S. export control laws, and may be subject to export or import regulations in other countries. If Ciena authorizes you to import or export the Software in writing, you shall obtain all necessary licenses or permits and comply with all applicable laws.

11. Limitation of Liability. ANY LIABILITY OF CIENA SHALL BE LIMITED IN THE AGGREGATE TO THE AMOUNTS PAID BY YOU TO CIENA OR ITS AUTHORIZED RESELLER FOR THE SOFTWARE. THIS LIMITATION APPLIES TO ALL CAUSES OF ACTION, INCLUDING WITHOUT LIMITATION BREACH OF CONTRACT, BREACH OF WARRANTY, NEGLIGENCE, STRICT LIABILITY, MISREPRESENTATION AND OTHER TORTS. THE LIMITATIONS OF LIABILITY DESCRIBED IN THIS SECTION ALSO APPLY TO ANY LICENSOR OF CIENA. NEITHER CIENA NOR ANY OF ITS LICENSORS SHALL BE LIABLE FOR ANY INJURY, LOSS OR DAMAGE, WHETHER INDIRECT, SPECIAL, INCIDENTAL OR CONSEQUENTIAL INCLUDING WITHOUT LIMITATION ANY LOST PROFITS, CONTRACTS, DATA OR PROGRAMS, AND THE COST OF RECOVERING SUCH DATA OR PROGRAMS, EVEN IF INFORMED OF THE POSSIBILITY OF SUCH DAMAGES IN ADVANCE.

12. General. Ciena may assign this License to an affiliate or to a purchaser of the intellectual property rights in the Software. You shall not assign or transfer this License or any rights hereunder, and any attempt to do so will be void. This License shall be governed by the laws of the State of New York without regard to conflict of laws provisions. The U.N. Convention on Contracts for the International Sale of Goods shall not apply hereto. This License constitutes the complete and exclusive agreement between the parties relating to the license for the Software and supersedes all proposals, communications, purchase orders, and prior agreements, verbal or written, between the parties. If any portion hereof is found to be void or unenforceable, the remaining provisions shall remain in full force and effect.

Publication history

January 2019

Standard Revision A

First standard release of this document for SAOS 6.18

Contents

About this document

This document describes how to configure the base system software on 39XX/51XX Service Delivery, Aggregation and Virtualization switches. This base system software is based on a common Service Aware Operating System (SAOS) code base designed to deliver consistent benefits across all Ethernet delivery, aggregation, and distribution configurations.

> ***Note:*** This base system software cannot be installed on any other Service Delivery Switches, Service Concentration Switches or Service Aggregation Switches.

This document provides information and examples for use in configuring base system software on any platform on which it is installed. It includes an explanation of the key features supported by the devices and provides example configurations for these features. Although these examples are useful in configuration, they are not meant to be used as a configuration template.

Conventions used in this document

Hyperlinks are indicated by blue text in this document.

In procedures, the following text conventions are used:

- `courier` text, for system responses
- *italic* text, for expected results
- **bold** text, for user input

Command syntax

A variety of symbols are used to indicate CLI command syntax. These symbols describe how to enter a command. They are not entered as part of the command itself. This table summarizes command syntax symbols.

Symbol	Description
< >	Encloses a variable or literal value that must be specified. Some examples include: server <IpAddress> priority <NUMBER: 1-7> dns <on|off> description <String[31]> For server <IpAddress>, the attribute can be entered as server 10.10.11.100 or server www.ciena.com. With priority <NUMBER: 1-7> the text within <> indicates that 1 - 7 are valid values. In the example of dns <on|off>, either the literal value of on or off is valid, such as dns on. For description <String[31]>, any string of up to 31 characters is entered.
{ }	Encloses a required value or list of required arguments. One or more values or arguments can be specified. For example, in the syntax: cfm mip create {vlan <VlanId>} {port <PortNameList>} [level <NUMBER: 0-7>] The vlan and port arguments are required. The level argument is optional.

|	Separates mutually exclusive items in a list, only one of which can be entered. For example, in the syntax:
	`dhcp client options set subnet <on|off>`
	Either on or off must be specified, for example:
	`dhcp client options set subnet on`

[ ]	Encloses an optional value or a list of optional arguments. One or more values or arguments can be specified. For example, in the syntax:
	`arp show [interface <Interface>]`
	You can enter a value for interface <Interface> or not. For example:
	`arp show`

Symbol	Description
{ [ ], [ ], [ ] }	Specifies a list of optional items where at least one must be specified.
…	Indicates the example has been abbreviated and that the actual display contains more information.
*	Indicates zero or more occurrences of what is preceding.

Documents in the 39XX/51XX documentation suite

For descriptions of documents in the 39XX/51XX Service Delivery, Aggregation and Virtualization Switches documentation suite, see *39XX/51XX Service Delivery, Aggregation and Virtualization Switches Product Fundamentals*.

New in this release

The following section summarizes documentation changes in *39XX/51XX Service Delivery, Aggregation and Virtualization Switches Base Configuration* for SAOS 6.18.

Port attributes

A couple of the administrative and operational attributes for ports are updated. See the Administrative and operational attributes for ports table on page 5-1 for details.

Link aggregation

Base link aggregation is now added because it no longer requires the Advanced Ethernet license. Multi-chassis link aggregation (MC-LAG) and the inter-chassis link (ICL) protocol remain in *39XX/51XX Service Delivery, Aggregation and Virtualization Switches Advanced Ethernet Configuration.*

Configuration fundamentals

This section provides an overview of the components and protocols required for configuring 39XX/51XX Service Delivery, Aggregation and Virtualization Switches:

- "Command line interface"
- "Command files and configuration files"
- "Ports"
- "Link Layer Discovery Protocol"
- "NETCONF/YANG"

Command line interface

39XX/51XX Service Delivery, Aggregation and Virtualization Switches are configured by means of a command line interface (CLI). Topics are:

- "User access" on page 2-1
- "Command schema" on page 2-2
- "Command syntax" on page 2-3

User access

The CLI hides commands depending on user access level and installed licenses. When entering a command at the prompt, ensure that you have the appropriate access level.

User access levels are

- super, for managing secure access to the switch through creation, deletion, and modification of user accounts
- admin, for making significant system state changes, modifying the system configuration, and performing execute commands
- limited, for system monitoring and gathering information about the configuration and performance of the system
- diag, for use when instructed by Ciena Customer Support to gather diagnostic information

This table describes the default names and passwords for accessing the CLI.

Table 2-1
Default user names and passwords

Group	User name	Password	Access rights
super	su	wwp	Read/Write/Create
admin	admin	wwp	Read/Write
limited	user	<empty>	Read-Only
diag	gss	pureethernet	Diagnostic

Note: Although the default limited user account cannot make configuration changes, the network operator must set a password for the account or delete the account.

For more information about user groups, refer to *39XX/51XX Service Delivery, Aggreagation and Virtualization Administration and Security*.

Command schema

In order to provide consistency across all the commands, all CLI commands follow a basic underlying schema or syntax:

```
<object> [<subobject>] <action> [instance] [<attributes>]
```

This table lists the elements of the command schema and provides a description of each element.

Table 2-2
Command schema

Element	Description
<object>	Identifies a feature or a basic object. On the surface, these seem like completely mismatched entities, however, if a device is considered to have an instance of each of these entities, then both features and the basic system objects can be considered as objects.
<subobject>	Subdivides a feature.
<action>	Describes the type of action to occur on the object or instance specified.
<instance>	Defines an recurrence of an object.
<attributes>	Identifies a pairing of a keyword and a value. A command can take multiple attributes and the attributes can be specified in any order relative to one another.

Command syntax

A variety of symbols are used to indicate CLI command syntax. These symbols describe how to enter a command. They are not entered as part of the command itself.

This table summarizes command syntax symbols.

Table 2-3
Command syntax symbols

Symbol	Description
< >	Encloses a variable or literal value that must be specified. Some examples include: server <IpAddress> priority <NUMBER: 1-7> dns <on|off> description <String[31]> For server <IpAddress>, the attribute can be entered as server 10.10.11.100 or server www.ciena.com. With priority <NUMBER: 1-7> the text within <> indicates that 1 - 7 are valid values. In the example of dns <on|off>, either the literal value of on or off is valid, such as dns on. For description <String[31]>, any string of up to 31 characters is entered.
{ }	Encloses a required value or list of required arguments. One or more values or arguments can be specified. For example, in the syntax: cfm mip create {vlan <VlanId>} {port <PortNameList>} [level <NUMBER: 0-7>] The vlan and port arguments are required. The level argument is optional.
|	Separates mutually exclusive items in a list, only one of which can be entered. For example, in the syntax: dhcp client options set subnet <on|off> either on or off must be specified, for example: dhcp client options set subnet on
[ ]	Encloses an optional value or a list of optional arguments. One or more values or arguments can be specified. For example, in the syntax: arp show [interface <Interface>] you can enter a value for interface <Interface> or not. For example: arp show

For more information about navigating the CLI, refer to *39XX/51XX Service Delivery, Aggregation and Virtualization Command Reference*.

Command files and configuration files

The command file directs the configuration of the SAOS devices. A configuration file is an ASCII file that contains a list of the CLI commands for that configuration.

Command file

The command file is the boot file for the SAOS platforms. The command file is an XML file that defines operations and associated attributes that can be automatically executed by the device.

Configuration file

The SAOS platforms can store multiple configuration files; however, only one configuration file can be active at a time. By default, configuration information is saved to a file called startup-config. The start-up config file is also the default load file. The parameters defined in the start-up-config file are applied when the device reboots unless an alternate file is specified. The current running configurations on SAOS devices are not saved to a configuration file unless specifically saved. This includes configuration changes made by the CLI, Simple Network Management Protocol (SNMP) or Network Configuration Protocol (NETCONF). If an SAOS device is rebooted without saving the configuration, all changes are lost.

Secure Zero-Touch Provisioning

Secure Zero-Touch Provisioning (SZTP) builds on Zero-Touch Provisioning (ZTP) by providing authenticated file access and encrypted file transmission. ZTP allows an SAOS platform to be provisioned automatically from a Dynamic Host Configuration Protocol (DHCP) server and a TFTP server. Once the SAOS platform is connected to the network and powered up, the device auto-configures according to instructions in the command file. SZTP is on by default.

Ports

Physical ports provide connectivity to other devices. Logical ports are created when multiple physical ports are joined in a Link Aggregation Group (LAG).

Physical ports provide connectivity to other devices, which is essential for any switching device. To aggregate bandwidth and provide link redundancy between two devices, physical ports are added to a Link Aggregation Group (LAG). The port management commands provide the ability to configure ports and troubleshoot connectivity.

Link Layer Discovery Protocol

Link Layer Discovery Protocol (LLDP) allows network equipment, for example, stations, switches, bridges, routers, to advertise their parameters for network topology discovery and management. Traditional network management protocols, such as SNMP, running at key locations, use layer 3 protocols to identify the devices connected to the network. The Link Layer Discovery Protocol is a layer 2 protocol, allowing precise discovery of the physical-link topology of the network. Devices act as LLDP agents, which drastically increases the network discovery performance of SNMP applications, as well as any system capable of accessing standard LLDP MIBs.

NETCONF/YANG

Network Configuration Protocol (NETCONF) provides a mechanism to install, manipulate and delete the configuration of a network device. YANG is the data modeling language that NETCONF uses. NETCONF uses XML-based data encoding for the configuration data and for the protocol message. NETCONF uses Remote Procedure Calls (RPCs) to communicate between the client and the server.

NETCONF provides monitoring and notification as part of network management. The NETCONF server usually resides on the network element that needs to be configured. The NETCONF client is typically a Network Management System that sends configuration information to the server. NETCONF can run 16 simultaneous sessions. See the NETCONF topology figure.

Figure 2-1
NETCONF topology

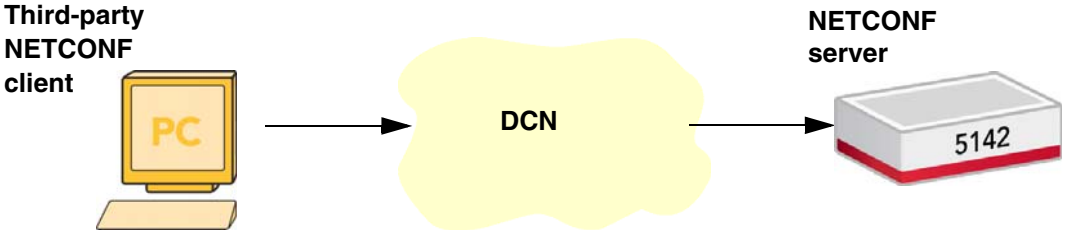

For more information, refer to “NETCONF/YANG configuration” on page 8-1.

Configuration management

This section provides explanations for implementing the basic set up of a device and describes the various options available to configure it, including:

This section provides the procedures for configuration management.

Accessing the CLI

To access CLI commands, connect to the device by establishing a Telnet session or through direct connection to the serial console port located on the front of the control module.

For related procedures, refer to *39xx/51xx Service Delivery, Aggregation and Virtualization Administration and Security*.

Configuration files

A device can store multiple device configuration files. However, only one configuration file can be active at a time. By default, configuration information is saved to a file called `startup-config`. The `startup-config` file is also the default load file. The parameters defined in the `startup-config` file are applied when the device reboots (unless an alternate file is specified). The current running configurations on a device are not saved to a configuration file unless specifically saved. This includes configuration changes made using the CLI or SNMP. If a device is rebooted without saving the configuration, all changes are lost.

This table lists the types of configuration.

Table 3-1
Types of configuration

Configuration	Description
Factory configuration	Any configuration made at the factory which cannot be changed by the network operator. When a Return to Factory Defaults (RTFD) operation is triggered, log files, and all configuration information created by the network operator are removed from the system. The system then uses the factory configuration when it boots up. ***Note:*** The ZTP mode of the 39XX/51XX Service Delivery, Aggregation and Virtualization Switches is not reset as part of an RTFD operation.
User configuration	Any configuration changes made to the default configuration by the network operator.
Default configuration	Hard-coded default settings found only in the source code, and not stored in a configuration file. The default configuration is subject to change from one release to the next.
Running configuration	The configuration that is running on the 39XX/51XX Service Delivery, Aggregation and Virtualization Switches. This configuration can be the startup configuration or can be a different configuration created from changes the network operator made to the 39XX/51XX Service Delivery, Aggregation and Virtualization Switches. Changes from the startup configuration are not necessarily saved to a configuration file unless the network operator saves them to the startup configuration. These changes can be lost if there is a system restart or if the failover does not support failing over to the running configuration. The network operator needs to be aware of unsaved configuration changes on the 39XX/51XX Service Delivery, Aggregation and Virtualization Switches.
Saved configuration	The network operator's configuration stored in a file.
Default load configuration	The configuration file that is selected by the network operator to be loaded when the 39XX/51XX Service Delivery, Aggregation and Virtualization Switches restarts.
Backup configuration	The configuration file selected by the network operator to be loaded at boot time if the primary saved configuration file cannot be loaded.
Revert configuration	The configuration file selected by the network operator to be loaded if the network operator fails to stop the revert timer before it expires.

Configuration flags

Configuration flags indicate that a change has been made in the configuration or that the configuration is busy. This table lists the configuration flags, how they are indicated, and what they mean.

Table 3-2
Configuration flags

Configuration flag	Symbol indication	Meaning
Config dirty	* 3930 *>	• Configuration has changed. It is different from the saved configuration. • The dirty flag disappears when the configuration is saved. The saved configuration is considered to be in ‘clean” status. • Asterisk appears after any actual configuration change or any potential change. If you set something to its current value, the command is allowed to mark the configuration as dirty. This behavior varies on a command-by-command basis.
Config busy	^ 3930 ^>	• The caret indicates a configuration is busy. A configuration can be busy if it is being augmented, printed, saved or some other blocking process is occurring. • The busy flag disappears when the configuration manager is no longer busy. The dirty flag may appear if the configuration is considered dirty.

Note: If you change a setting and change it back to its prior value, the configuration does not have a clean status. You can verify whether the running configuration matches the saved configuration by running the `configuration show differences-from-saved` command.

Events and traps notifications are generated when configuration changes are made on a device. This table describes the events and traps.

Table 3-3
Events and Traps

Event and Trap	Description
configSave trap	• Configuration is saved on the device. • Contains the time of the configuration save and the time of the last configuration save. • Configuration changes made on the device by means of the CLI, SNMP or NETCONF.
configSave event	• Configuration changes are saved on a device. • Contains the time of the configuration save and the time of the last configuration change. • Configuration changes made on the device by means of DHCP or DHCPv6.
configChange trap	• Can be enabled/disabled by means of the CLI. • Configuration changes are made on the device. • The number of traps generated is limited. Traps cannot be bundled for different users. • ConfigChange notifications are not generated until the system is operationally enabled. Contains the following information: • Context of the change: by means of the CLI, SNMP or NETCONF • User/IP address from where the change request originated from • Timestamp of the configuration change. • Subsystem where the configuration change occurred. • Context of the create/modify/delete change. • To/from data change

Configuration procedures are:

- "Saving configuration changes" on page 3-13
- "Checking the syntax of commands in configuration files" on page 3-15
- "Adding commands to the running configuration" on page 3-17
- "Displaying configuration files" on page 3-20
- "Activating alternate configurations" on page 3-22
- "Displaying the default configuration files" on page 3-23

Secure Zero Touch Provisioning

Secure Zero Touch Provisioning (SZTP) builds on Zero Touch Provisioning (ZTP) by providing authenticated file access and encrypted file transmission. SZTP file transfers are done using SFTP. DHCP option 125 contains the Ciena Bootfile URL which specifies the file transfer protocol, the file server address, and the name of the bootfile, otherwise known as the XML command file.

An Advanced Security license must be pre-installed by Ciena as it is required to use SFTP. SFTP credentials must also be pre-installed so that the SFTP server can grant access. SFTP credentials consist of a username and password which are stored in non-volatile/flash/persistent storage. The username and password are encrypted in this storage. The username and password are not overwritten, cleared or reset by returning to factory default. They can be set or unset by the user using the CLI commands.

The SZTP overview figure shows an overview of SZTP. It is identical to ZTP with the exception that server/protocol information is obtained from a specific DHCP option using SFTP instead of TFTP.

Figure 3-1
SZTP overview

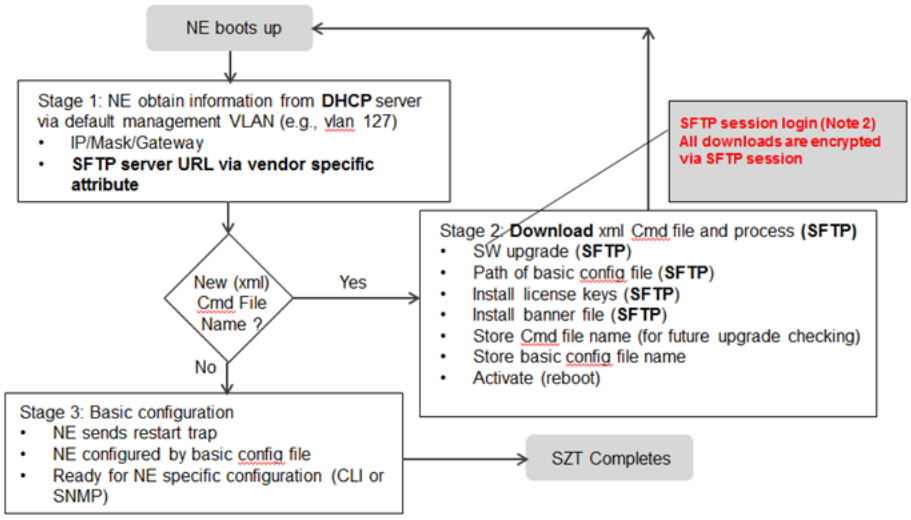

SAOS supports DHCP option 125 for vendor-identifying vendor-specific information and the Ciena Command File URL sub-option (0x10). The Command File URL is parsed into a scheme (protocol), a hostname, and a path/file name. The scheme identifies how the command file can be accessed. It is used as the default scheme for accessing any files referenced by the command file unless otherwise overridden. SAOS 6.13 and later support the URL prefix “sftp://” for SZTP. The server identifies the server from which to obtain the command file. It is similar to the TFTP server name (option 66) in DHCP. Unless it is overridden, this server is used to access the files referenced by the command file. The path/file name is used to specify the specific path/filename of the command file. It is handled like the bootfile name (option 67) in DHCP.

The Ciena Command File URL sub-option takes precedence over options 66 and 67 if it is enabled, is sent by the server and contains valid syntax and either

- the scheme is SFTP or
- the scheme is SFTP, and a security license and sftp credentials are present

In any other case, the Ciena Command File URL is ignored and options 66 and 67 are used.

SZTP procedures are:

- “Configuring Secure Zero Touch Provisioning username and password” on page 3-29
- “Displaying Secure Zero Touch Provisioning information” on page 3-30

TPID rotation during zero touch provisioning

When a software delivery switch boots, it checks for the presence of the start-up configuration file. If this file is not present, TPID rotation is enabled.

If TPID rotation is enabled when the DHCPv4 or DHCPv6 client is started, a 90-second TPID rotation timer starts. When the timer expires, any operating DHCPv4 or DHCPv6 client stops, the TPID of egress frames from the remote interface VLAN (127 by default) is changed to the next value, and any operating DHCPv4 or DHCPv6 client is re-started.

The TPID values cycle through this list, with the last value followed by the first value in a continuous rotation:

- 0x8100
- 0x88a8
- 0x9100
- Untagged

When the DHCPv4 or DHCPv6 client is started, it attempts to contact a DHCPv4 or DHCPv6 server and obtain a lease. Until it receives a lease, the client sends a DHCP discover message or DHCPv6 solicit message about every 30 seconds (discovery interval). If it receives a valid lease, the client is satisfied and stops sending DHCPv4 or DHCPv6 messages until the lease renewal timer (t1) expires, and the TPID rotation timer stops.

The DHCP client sends a discover message shortly after it starts. If TPID rotation is not enabled and no response is received from a server, the client continues to send discover messages about every 30 seconds. If TPID rotation is enabled, the sequence of discover messages is interrupted every 90 seconds by a change of TPIDs. The DHCP client is stopped and re-started and given a chance to find a server using the next TPID in the sequence.

In this example, if there are no DHCP servers on the network to which the device's remote interface is connected and there is also no default configuration file on the device, the device sends DHCP messages as follows:

- Send a Discover message with TPID 0x8100.
  Wait 30 seconds.
  Send a Discover message with TPID 0x8100.
  Wait 30 seconds.
  Send a Discover message with TPID 0x8100.
  Wait 30 seconds.
  (The TPID rotation timer expires.)
- Send a Discover message with TPID 0x88a8.
  Wait 30 seconds.
  Send a Discover message with TPID 0x88a8.
  Wait 30 seconds.
  Send a Discover message with TPID 0x88a8.
  Wait 30 seconds.
  (The TPID rotation timer expires.)
- Send a Discover message with TPID 0x9100.
  Wait 30 seconds.
  Send a Discover message with TPID 0x9100.
  Wait 30 seconds.
  Send a Discover message with TPID 0x9100.
  Wait 30 seconds.
  (The TPID rotation timer expires.)
- Send a Discover message untagged.
  Wait 30 seconds.
  Send a Discover message untagged.
  Wait 30 seconds.
  Send a Discover message untagged.
  Wait 30 seconds.
  (The TPID rotation timer expires.)

- Repeat from the top.

If the user changes any attribute of the remote interface or the DHCP or DHCPv6 client while the TPID rotation is on, the TPID rotation stops. If TPID rotation makes any changes to the remote interface, such as the TPID rotation timer has expired at least once, those changes are reverted and the remote interface settings are returned to their default values.

TPID rotation behavior

This table shows the TPID rotation states and how they occur.

Table 3-4
TPID rotation states

TPID rotation state	Description
Off	The TPID rotation state is initially “off”. The TPID rotation state is set to “off” and the frame type restored to TPID 0x8100 when the user sets any attribute of the remote interface of the DHCP or DHCP v6 client. The disabled TPID rotation state occurs for one of these reasons: • A DHCPACK message or DHCPv6 relay message is received by the DHCP or DHCPv6 client. • The user sets any attribute of the remote interface or of the DHCP or DHCPv6 client. When the user makes these attribute changes and TPID rotation has been running for at least 90 seconds so that it has changed the TPID at least once, those changes made automatically by TPID rotation revert to their initial, default states, including using TPID 0x8100.
On	The device boots and determines that there is no default load file or backup load file on the device. The TPID rotation timer can only expire while the TPID rotation state is “on”.

The system can perform zero touch provisioning over a network configured to exchange frames with the remote interface using any of these frame types:

- 0x8100

- 0x88a8
- 0x9100
- Untagged

This also includes obtaining an IP address, an XML command file and any files specified in the command file, using TFTP, FTP and SFTP, and obtaining the command file name and location from either options 66 or 67 or option 125.

Device configuration during TPID rotation

On a defaulted device, the TPID rotation controller configures administrative parameters of a remote management VLAN and its members as per the TPID value in cycle until DHCPv4/DHCPv6 hunting is successful.

If the dataplane configuration required for connectivity of the device to a DHCPv4/DHCPv6 server is known, then see "Manually configuring a device for IPv4/IPv6 leasing" on page 3-33.

Caveats

During TPID rotation, these administrative parameters of remote management VLAN and its member ports change:

- egress-tpid of remote management VLAN to 0x8100/0x88a8/0x9100
  - `vlan set vlan <remote_management_vlan> egress-tpid <8100|88a8|9100>`
- vlan-ethertype-policy of ports which are members of the remote management VLAN to inherit vlan-tpid
  - `virtual-circuit ethernet set port <port> vlan-ethertype-policy vlan-tpid`
- Port's attributes pvid and egress-untagged-vlan equivalent to remote management VLAN:
  - `port set port <port> pvid <remote_management_vlan> egress-untag-vlan egress-untag-vlan <remote_management_vlan`

Note 1: These changes are seen as part of the configuration show command. The configuration is not marked as dirty. No syslogs are generated for the configuration change due to TPID rotation.

Note 2: Configuring the port's attribute "acceptable-frame-type" to "all" to accept all packets opens the channel for all untagged packets.

Note 3: If the IPv4/IPv6 address is leased for the device while TPID rotation is active, changing these attributes of members of remote management VLAN is not recommended and may impact IPv4/IPv6 leasing or connectivity to the device: acceptable-frame type, pvid, egress-untag-vlan, and vlan-ethertype-policy.

Note 4: If the IPv4/IPv6 address is leased for the device during TPID rotation or if TPID rotation is active, changing the egress-tpid of the remote management VLAN is not recommended and may impact IPv4/IPv6 leasing or connectivity to the device.

Note 5: If ports are added to the remote management VLAN when TPID rotation is active, then their administrative parameters are modified as per active TPID value when TPID rotation switches to the next value in the cycle.

Note 6: If ports are removed from the remote management VLAN when TPID rotation is active then the administrative parameters remain the same as modified during the TPID rotation.

Note 7: During “untagged” mode of TPID rotation, DHCPv6 hunting is not supported. If the DHCPv6 server is reachable by an untagged network, then the attribute of port (member of remote management VLAN) “acceptable-frame-type” must be configured to “all” to accept untagged and tagged packets during the “untagged” mode of the TPID rotation.1

DHCP client re-initiation on link transition

This user-configurable feature triggers DHCP client re-initiation when a management link goes down. The management link can be either a local or remote NNI link. The DHCP client re-initiates and sends a discover again to get new IP configuration if any NNI fault occurs. DHCP client re-initiation occurs automatically, without human intervention. This feature applies to Service Delivery Switches running SAOS 6.x software.

Usually the DHCP server is configured to provide the required IP address, subnet and gateway and this information is valid for the lease time (in seconds) returned by the DHCP server. The least must be renewed before it expires; otherwise the configuration is no longer valid and should not be used by the system. Along with the IP address configuration, the DHCP server can send a list of DHCP or BOOTP options that the system uses to configure various components of the system.

This feature modifies the default behavior of DHCP. When it is enabled, DHCP IP, lease and configuration information are re-initialized if any management link fault occurs. For remote management, NNI link faults can be caused by any of these scenarios:

- Physical layer NNI fault
- Administrative NNI port disable
- Re-configuration of the NNI port in a device that results in the port unable to forward traffic

For example, if two NNI on a switch are connected with different operators of a provider network, this feature allows the switch (and customer CPE) to obtain new DHCP IP configuration information from a backup NNI if the active NNI goes down due to any link fault.

A management interface provides IP connectivity to the system for management purposes. The DHCP client re-initiates and refreshes all IP configuration information received from the DHCP server when this feature is enabled. Impacts of this feature over management interface IP configuration provided by the system are:

- Local management interface—If the DHCP client interface is set to local, the DHCP client re-initiates in case of any local management fault conditions.
- Remote management interface—This in-band management interface provides IP connectivity through one or more user ports on the system. The remote management interface can be configured using:
  - — Management VLAN: The DHCP client re-initiates if the active NNI port of the management VLAN goes down or any possible re-configuration of the NNI port that results in the VLAN unable to forward management traffic. The active NNI port is a member of the remote management VLAN from which the management IP address was learned. In the case of multiple NNI ports when the active and forwarding NNI port is faulted, the DHCP client re-initializes and gets the IP configuration from the other active NNI port.
  - — Management VS: The remote management interface can be associated with the Management VS instead of the Management VLAN. Only VPLS VS can be associated with the remote management interface. This feature is not supported for Management VS fault. The DHCP client does not re-initialize in case of any fault in the Management VS due to failure or re-configuration.

The following actions over local or remote management interface cause a DHCP client to re-initiate when DHCP is operationally enabled and the release on link transition feature is enabled:

- Local or remote management interface is disabled due to a physical link fault
- Remote management VLAN has changed
- Active NNI port member of the remote management VLAN is disabled
- Active NNI port is removed from the remote management VLAN
- Active NNI port is also a member of a link aggregation group and the aggregation port is disabled

The DHCP re-initiation on link transition feature is enabled and disabled through the CLI. See Procedure 3-16, “Activating DHCP client re-initiation on link transition” on page 3-35 for details. The DHCP YANG model has also been enhanced to include objects that support feature on/off.

Procedure 3-1
Saving configuration changes

To permanently save configuration changes, you must save the running configuration to a configuration file. To save the running configuration, use the `configuration save` command. By default, the command saves the current configuration to the default configuration file, `startup-config`.

A configuration save event is generated when the configuration is saved on the device. This event contains the time of the configuration save and the last configuration change. The configuration change event is generated when a configuration change is made on the device by means of the CLI, SNMP, NETCONF, DHCP, DHCPv6. The change event is not generated until the system is in an operationally enabled state.

A configuration trap is generated when the configuration is saved on the device. The trap contains the time of the configuration save and the time of the last configuration change. The generation of the trap can be enabled or disabled from the CLI. It is enabled by default. A configuration trap contains the

- context in which the change was done (CLI, SNMP, NETCONF)
- IP address from where the change request originated from
- timestamp of the configuration change

By saving alternate versions of command files, you can store multiple configuration files for running different configurations of the system. For example, you can save configuration to an alternate file as a backup to restore to a previous configuration or you can store a configuration file for configuring another device of the same family.

> ***Note 1:*** An asterisk * indicates that a configuration is dirty. This means that the configuration needs to be saved. When the configuration is saved, the asterisk disappears.
>
> ***Note 2:*** A caret ^ indicates that the configuration is busy. This means that a configuration change is being saved, augmented, printed or some other blocking process is occurring. When the caret disappears, the configuration may be flagged as dirty.

Step	Action

To save the running configuration

1 Save the running configuration:

```
configuration save
```

To save to an alternate filename

2 Save the configuration to an alternate filename:

```
configuration save filename <filename>
```

where

filename <filename> is the configuration file name.

—end—

Example

This example saves the default configuration file.

```
configuration save
```

This example saves to an alternate filename.

```
configuration save filename myConfig
```

Procedure 3-2
Checking the syntax of commands in configuration files

Check the syntax of commands before:

- adding the commands to the default configuration from a file
- activating an alternate configuration file

You can check a configuration file in the file system or on an xFTP server. Also, you can have each command displayed on the screen as the check occurs or hide the progress of the file check.

Step	Action
1	Check the syntax of commands in a configuration file:

```
configuration check {default-server|default-ftp-
server|default-tftp-server|default-sftp-server|
{tftp-server <ip-host-str> [server-port <INTEGER:
1..65535>]}|
{ftp-server <ip-host-str> [login-id <username>
[<password-attr>|<encrypted-password-attr>|<echoless-
password-attr>][server-port <INTEGER: 1..65535>]}|
{sftp-server <ip-host-str> login-id <username>
{<password-attr>| <encrypted-password-attr>|<echoless-
password-attr>}[server-port <INTEGER: 1..65535>]}}
[verbose] [hide-progress]
```

where

default-server	indicates to use the default xFTP server.
default-ftp-server	indicates to use the default FTP server.
default-tftp-server	indicates to use the default TFTP server.
default-sftp-server	indicates to use the default SFTP server.
sftp-server <IP address or host name>	is the SFTP server.
tftp-server <IP address or host name>	is the TFTP server.

where

server-port <INTEGER: 1..65535>	is the server-port number.
ftp-server <IP-host-str>	is the FTP server name.
login-id <username>	is the FTP/SFTP username.
password-attr>	enters the password in clear text.
encrypted-password-attr	sets the password using a pre-encrypted string.
echoless-password-attr	engages an echoless password collector.
server-port <INTEGER: 1..65535>	is the server-port to connect to.
verbose	prints commands before parsing.
hide-progress	hides the progress of the parse.

—end—

Example

This example shows sample output for a file without syntax errors.

```
> configuration check filename startup-config
ConfigMgr: Checking configuration file /flash0/config/startup-config

Checking 36265-line configuration file /flash0/config/startup-config...
ConfigMgr: Finished checking configuration file /flash0/config/startup-config
```

This example shows sample output for a file with syntax errors.

```
> configuration check filename myConfigAdd.txt tftp-server 192.168.41.68
ConfigMgr: Checking configuration file /ram/temp30

Checking 1-line configuration file /ram/temp30...
CONFIG ERROR: Config file line 1:
      virtual-switch private-forwarding-groups disable 1000
ConfigMgr: Finished checking configuration file /ram/temp30
ERROR: Errors were found in the configuration file
```

Procedure 3-3
Adding commands to the running configuration

Add commands to the running configuration from a file when there are changes in network topology, for example, when adding a service, such as quality of service.

CAUTION
Risk of service impact
Ensure that commands added to the running configuration do not impact existing services. For example, changing the G.8032 port does impact existing services.

By default, the commands are added from a specified file stored on the file system at /flash0/config.

Note 1: If you want to apply the added commands to the default startup command file, save the configuration.

Note 2: Not all CLI commands are allowed within an augmentation file. For example

- all variations of “configuration save”
- configuration set default-save-filename
- configuration set default-load filename

You get a SHELL PARSER FAILURE if you attempt to use a disallowed command.

Step	Action

To add configuration commands to the running configuration

1 Add configuration commands to the running configuration by means of a file:

```
configuration augment {default-server|default-ftp-
server|default-tftp-server|default-sftp-server|
{tftp-server <ip-host-str> [server-port <INTEGER:
1..65535>]}|
{ftp-server <ip-host-str> [login-id <username>
[<password-attr>|<encrypted-password-attr>|<echoless-
password-attr>][server-port <INTEGER: 1..65535>]}|
{sftp-server <ip-host-str> login-id <username>
{<password-attr>| <encrypted-password-attr>|<echoless-
password-attr>}[server-port <INTEGER: 1..65535>]}}
{filename <String>} [verbose] [hide-progress]
```

where

default-server	indicates to use the default xFTP server.
default-ftp-server	indicates to use the default FTP server.
default-tftp-server	indicates to use the default TFTP server.
default-sftp-server	indicates to use the default SFTP server.
sftp-server <IP address or host name>	is the SFTP server.
tftp-server <IP address or host name>	is the TFTP server.
ftp-server <IP address or host name>	is the FTP server.
server-port <INTEGER: 1...655535>	is the server-port number.
ftp-server <IP-host-str>	is the FTP server name.
login-id <username>	is the FTP/SFTP username.
password-attr>	enters the password in clear text.
encrypted-password-attr>	sets the password using a pre-encrypted string.

where

echoless-password-attr	engages an echoless password collector.
server-port <INTEGER: 1...655535>	is the server-port to connect to.
filename <String>	is the configuration file name.
verbose	prints commands before parsing.
hide-progress	hides the progress of the parse.

To stop a configuration augment in progress from the same session

2 Press Ctrl+C.

3 Stop the configuration augment in progress:

```
configuration stop
```

Note: Commands that have been added to the running configuration are not removed.

To stop a configuration augment in progress from a different session

4 Stop the configuration augment in progress:

```
configuration stop
```

—end—

Procedure 3-4
Displaying configuration files

Display configuration files.

Step	Action

1 Display the configuration:

```
configuration show [status] [line-numbered] [differences-
from-saved] [include-default-settings] [brief]
```

where

status	displays the current state of the configuration system.
line-numbered	displays with line numbers.
differences-from-saved	displays changes since last saved.
include-default-settings	displays the current configuration including default settings.
brief	displays the current running configuration.

2 Search the configuration file:

```
configuration search file <String>] [lines <NUMBER: 0-
40>] {string <String>}
```

where

file <String>]	is the filename (no path).
lines <NUMBER: 0-40>	is the search window size.
string <String>	is the search string.

—end—

Example

This example displays the differences between the running configuration and the saved configuration file.

```
> configuration show differences-from-saved
diff /flash0/config/Mcast_Aggs_DG /ram/65700.out
4,6c4,6
< ! Created:       Mon May  5 14:04:02 2008
---
> ! Created:       Mon May  5 14:09:04 2008
```

This example displays sections of a configuration file containing a specific string.

```
> configuration search string dhcp
dhcp client disable
dhcp client set interface local
```

Procedure 3-5
Activating alternate configurations

Activate an alternate configuration to change the entire configuration, for example, to roll back to a previous configuration if an updated configuration does not meet expectations. Alternate configurations are typically activated in a lab environment.

You can

- activate a saved alternate command file. The system runs a basic syntax check, restarts the system, and then loads the specified file.
- restore the configuration while leaving the file system intact: license keys, logs or configuration files are not removed. The system renames the startup-config file to startup-config.bak so it is not loaded.

After running any of the commands to activate an alternate configuration, you have 10 seconds to cancel. To cancel, press Ctrl+C.

Step	Action

To activate an alternate command file

1 Activate an alternate command file:

```
configuration reset-to-user-config filename <filename>
```

where

filename <filename> is the configuration file name

To restore the configuration while leaving the file system intact

2 Restore the configuration while leaving the file system intact:

```
configuration reset-to-defaults
```

The system restarts without loading a configuration file.

—end—

Example

This example activates an alternate command file named myConfig.

```
>configuration reset-to-user-config filename myConfig
WARNING:   You cannot abort the reset operation once it has started.
Proceeding in 10 seconds.  Press <CTRL>C to abort...--break--
*>
```

Procedure 3-6
Displaying the default configuration files

The default configuration files are

- save
- load

Step	Action
1	Display the default save, load, and backup load files: `configuration list`

—end—

Example

This example shows sample output for the configuration list command.

```
> configuration list

+------------------------------------------------------------------------------+
|      Configuration Files                                                     |
+------------------------------------------------------------------------------+
| startup-config                                                               |
| test                                                                         |
+------------------------------------------------------------------------------+
| Default Save File: test                                                      |
| Default Load File: test                                                      |
+------------------------------------------------------------------------------+
```

Procedure 3-7
Restoring default configuration to user defaults

You can

- reset configuration to user defaults

After running any of the commands to activate an alternative configuration, you have ten seconds to cancel. To cancel, press Ctrl+C.

Step	Action
1	Reset configuration to user defaults: `configuration reset-to-user-config [filename <filename>]` where filename <filename> is the configuration file name.

—end—

Example

This example resets the configuration to user defaults.

```
configuration reset-to-user-config filename myConfig
```

Procedure 3-8
Resetting to factory default configuration

You can reset the device to factory configuration defaults by using the CLI, by means of SNMP, or by pressing the reset button on the device.

CAUTION
Loss of configuration information
A factory configuration comprises the default settings in place when the system was shipped.

When you reset a device to its factory default settings, all configuration and file system changes are lost, including saved configuration files and log files.

A factory reset sets the system back to default conditions and removes all configuration and log files. The system shuts down and reboots.

Factory default settings vary depending on the platform and port number. Refer to *Product Fundamentals* for a listing of default settings applicable to all platforms and ports and default settings that vary by platform and port.

These customer-visible operational items are explicitly cleared during a factory reset:

- Any configuration changes made by any and all users
- External configuration parameters stored outside of configuration files:
  - the configuration load file name
  - the configuration save file name
  - debug settings
- Event, alarm and performance history logs
- Last command file downloaded when DHCP makes a connection either through ZTP or automatically if DHCP is enabled without ZTP
- All configuration files stored in /flash0/config
- All settings of interfaces (IP/subnet)
- All SSH and MD5 key files

These customer-visible operational items are explicitly preserved during a factory reset:

- All software packages installed (including the running package)
- All Software License keys

- Cumulative system uptime — is the sum of the archived system uptime (sampled at SAOS start-up) and the time that the current SAOS server instance has been running.
- Archived system uptime — the uptime currently saved in the device. It differs from cumulative uptime as it is only updated once a day.
- Highest temperature
- Number of High Temperature Violations
- Lowest temperature
- Number of Low Temperature Violations
- Number of Soft Resets
- Total number of Resets
- Last Reset Reason
- Total number of Guardian Reboots — a reboot triggered by the Guardian software watchdog.
- Number of consecutive Guardian Reboots

Step	Action

To restore the device to factory defaults

1 Restore the device to factory defaults:

```
configuration reset-to-factory-defaults
```

The system shuts down and reboots.

—end—

Procedure 3-9
Setting the default configuration files

When a 39XX/51XX switch is restarted, it loads configuration from the default load file. The default saved configuration file is the target of the network operator's configuration save operation: normally these are the same file.

You can choose to save settings to an alternate configuration file when a configuration save operation is performed, for example, if you are incrementally building a new configuration that is not yet ready to take the place of the default load file. Once the new file is ready to be deployed, change the default load file to also point to the new file.

You can

- set the default file for saving configuration
- set the default file for loading configuration

Step	Action

To set the default file for saving configuration

1 Set the default file for saving configuration:

```
configuration set default-save-filename <FileName>
```

To set the default file for loading configuration

2 Set the default file for loading configuration:

```
configuration set default-load-filename <FileName>
```

—end—

Procedure 3-10
Resetting default configuration files to factory default files

You can

- reset the default file for saving configuration
- reset the default file for loading configuration

Step	Action

To reset the default file for saving configuration

1 Reset the default file for saving configuration:

```
configuration unset default-save-filename
```

To reset the default file for loading configuration

2 Reset the default file for loading configuration:

```
configuration unset default-load-filename
```

—end—

Procedure 3-11
Configuring Secure Zero Touch Provisioning username and password

You can configure the Secure Zero Touch Provisioning (SZTP) server account username and password.

Step	Action

To set the SZTP username and password

1 Configure the SZTP server account username:

```
dhcp security set {username <USERNAME[1..32]> [echoless-password]}
```

where

username <USERNAME [1..32]>	is the SFTP server account user name.
echoless-password	engages an echoless password collector

To set the SZTP attributes to default values

2 Set the SZTP attributes to default values:

```
dhcp security unset {[username][password]}
```

—end—

Procedure 3-12
Displaying Secure Zero Touch Provisioning information

You can display SZTP information.

Step	Action
1	Display SZTP information: `dhcp security show`

—end—

Example

This example shows the command output from the dhcp security show command:

```
dhcp security show

+-- DHCP BOOTFILE TRANSFER SECURITY CREDENTIALS --+
| Credential                          | State    |
+-------------------------------------+----------+
| SFTP username                       | Unset    |
| SFTP password                       | Unset    |
+-------------------------------------+----------+
```

Procedure 3-13
Setting the TPID rotation state to off

You can only set dhcp tpid-rotation state to off. It cannot be set to on. TPID rotation only comes on when the device is booted with no default configuration load present.

Step	Action
1	Set the TPID rotation state to off: `dhcp tpid-rotation set off`

—end—

Procedure 3-14
Displaying TPID rotation

Display TPID rotation parameters.

Step	Action
1	Display TPID rotation: `dhcp tpid-rotation show`

—end—

Example

This example displays output from the dhcp tpid-rotation show command.

```
+-------- TPID ROTATION STATE --------+
| Parameter                | Value    |
+--------------------------+----------+
| Rotation State           | On       |
| Last Tx Rotation TPID    | 88A8     |
+--------------------------+----------+
```

Procedure 3-15
Manually configuring a device for IPv4/IPv6 leasing

In the case where you know the configuration required for connectivity of the device to a DHCPv4/v6 server, then you can manually configure a device for IPv4/IPv6 leasing.

Step	Action

To manually configure a device for IPv4/IPv6 leasing when the DHCPv4/DHCPv6 server is reachable by a tagged network

1 Set “egress-tpid” of the remote management VLAN:

```
vlan set vlan <remote_management_vlan> egress-tpid
<8100|88a8|9100>
```

where

vlan <remote_management_vlan>	is the remote management VLAN name.
egress-tpid <8100|88a8|9100>	sets the egress-tpid.

2 Set the attributes of the port which is a member of the remote management VLAN and that is providing connectivity to the DHCPv4/DPCPv6 server:

```
virtual-circuit ethernet set port <port> vlan-ethertype-
policy vlan-tpid
```

where

port <port>	is the port which is a member of the remote management VLAN and that is providing connectivity to the server.

To manually configure a device for IPv4/IPv6 leasing when the DHCPv4/DHCPv6 server is reachable by an untagged network

3 Set the port’s attributes which are equivalent to the remote management VLAN:

```
port set port <port> pvid <remote_management_vlan>
egress-untag-vlan <remote_management_vlan>
```

where

pvid <remote_management _VLAN>	is the remote management PVID VLAN name.
egress-untag-vlan <remote_management _VLAN>	is the egress-untagged remote management VLAN.

4 Set the port's acceptable-frame-type attribute to accept untagged and tagged packets:

```
port set port <port> acceptable-frame-type all
```

—end—

Procedure 3-16
Activating DHCP client re-initiation on link transition

Use this procedure to activate DHCP client re-initiation in case a management link goes down. This setting is disabled by default.

Step	Action

To enable this feature

1 Turn on DHCP re-initiation on link transition:

```
> dhcp client set release-on-link-transition on
```

To disable this feature

2 Turn off DHCP re-initiation on link transition:

```
> dhcp client set release-on-link-transition off
```

—end—

Management interface configuration

The preferred source interface for management traffic feature gives users the option to use a configured Internet Protocol (IP) interface (including loopback) address as a source address for the management traffic in a network.

One VLAN is reserved to create a simulated loopback IP. The reserved VLAN must contain all the ports that have management EMS/NMS reachability.

One IP interface is created with a /32 mask and the reserved VLAN is assigned to it. This IP interface is used for all the management protocols. The subnet the IP interface is part of can be advertised into Border Gateway Protocol (BGP), for example, as used in seamless MPLS. See the “Seamless MPLS configuration” chapter in *39XX/51XX Service Delivery, Aggregation and Virtualization Switches MPLS Configuration* for details on this usage.

All the management protocols (listed in the introduction to “Setting a preferred source IP” on page 4-4) bind the source IP (IP of an already created IP interface) to the connect socket. If the IP is binded to the socket, the kernel does not alter this IP with the outgoing interface.

This feature is supported on the 390x, 3926m, and 3928 platforms.

Scope

Currently only IPv4 support for management protocols has been planned. IPv6 support cannot be done because of the Border Gateway Protocol (BGP) limitation (BGP does not support IPv6).

By default, this feature is disabled for all the management protocols.

Only traffic initiated by the node is handled. There is no change in the IP header for sessions not initiated from the node (in case of response).

Prerequisites

A unique IP interface only for management traffic must be configured. This IP interface is used when enabling this feature for any protocol. Management traffic uses the IP configured for this interface as the source IP in the IP header.

BGP distributes the routes for this IP interface in the complete network to handle the response packets.

BGP policy makes sure no management route is learned by the control EMS and no control route is learned by the management EMS.

Operational flow

Using the “set” CLI enables this feature, ensuring protocol traffic uses the already configured IP interface as a preferred source address. The IP interface address is used in the IP header as the source IP. The actual outgoing IP interface to send a packet out is based on routing.

Using the “unset” CLI disables this feature, causing the management traffic to choose the IP of the outgoing interface as its source IP.

IP interface creation and deletion is managed by the IP manager. But this IP interface is dependent on different management protocols. So if this IP interface is referred to by any of the protocols, the user cannot delete it. To delete the IP interface, the user needs to remove associations with this interface from all the management protocols.

Limitations

Feature limitations include:

- Configuration of the preferred source IP through SNMP and through NETCONF YANG is not supported.
- These protocols have not been addressed:
  - Dying gasp
  - ZTP
  - DHCP client
  - Netconf
  - OCSP
- There is no separate CLI for RMON to enable the feature. RMON is dependent on xFTP configurations.
- There is no separate CLI for Dot1x to enable the feature. Dot1x is dependent on RADIUS configurations to configure the preferred source IP address.

- FTP, TFTP, and SFTP use common CLI to enable and disable this feature.
- For Telnet client implementation, the “socat” utility has been used. Socat performs IAC negotiation while establishing the session. Users see some special characters in the login banner due to IAC negotiation. There is no impact on Telnet client functionality.
- For GET/PUT operations using SFTP, setting the preferred source IP for SSH along with xFTP is required.
- The IP interface is configured to simulate another loopback interface for management traffic. An unnamed IP interface must not be configured to achieve preferred source IP functionality.

Procedure 4-1
Setting a preferred source IP

Use this procedure to set a preferred source IP for a management interface. This procedure can also be used to unset a preferred source IP.

Management protocols for which a preferred source IP can be set and unset include:

- Telnet client
- SSH client
- FTP client
- TFTP client
- SFTP client
- RADIUS client
- RadSec client
- TACACS client
- SNMP trap
- Syslog client
- Syslog TLS client
- NTP client
- RMON
- DNS client
- Dot1x

Step	Action

Set a preferred source IP

1 Set a preferred source IP:

```
<protocol> <set> preferred-source-ip <interface name>
```

where

<protocol>	is one of the management protocols listed in the introduction to this procedure.
<set>	is the action being performed.
<interface name>	is the name of the already configured management interface you want as the preferred source IP.

Unset a preferred source IP

2 Unset a preferred source IP:

```
<protocol> <unset> preferred-source-ip
```

—end—

Examples

```
> telnet client set preferred-source-ip <interface name>

> telnet client unset preferred-source-ip
```

Port management

This section explains how to configure physical and logical port attributes.

Physical ports provide connectivity to other devices, which is essential for any switching device. To aggregate bandwidth and provide link redundancy between two devices, physical ports are added to a Link Aggregation Group (LAG). The port management commands provide the ability to configure ports and troubleshoot connectivity.

Port management addresses:

- "Port attributes"
- "Port statistics"
- "Transceivers"

This section provides the procedures for port management.

Port attributes

This table describes administrative and operational attributes for ports.

Table 5-1
Administrative and operational attributes for ports

Attribute	Description
General	
port <port>	A 32 character string representing the name of the physical port or LAG. For physical ports, the name represents the port's physical location identifying the chassis module and port in the format: <ModuleNumber>.<PortNumber> Whenever a platform has a single module, the component number is left out of the name. For example, the 3930 platform is single module, so each physical port is named only with the port number (1 through 10). The 5150 platform supports multiple modules, so each port on the second and third module is named with the module and port numbers (2.1, 2.2, 3.1, 3.2).
link-flap-detect <on | off>	Configuration for monitoring link state transitions.

Table 5-1
Administrative and operational attributes for ports

Attribute	Description
link-flap-count <NUMBER: 1-64>	Sets the link transition count that causes a link flap event. Default is 5.
link-flap-detect-time <SECONDS: 1-600>	Sets the link flap detection window in seconds. Default is 10.
link-flap-hold-time <SECONDS: 0-600>	Sets the link flap port re-enable time in seconds. Default is 300.
hold-off <on | off>	Sets the port hold-off state. Default is off.
hold-off-time <NUMBER: 3-20>	Sets the port hold-off time in deciseconds (1 ds = one-tenth of a second or 100 ms). Default is 10 ds.
acceptable frame type <all | tagged-only | untagged-only>	Designates the treatment of received frames to allow all, tagged-only, or untagged-only.
advertised-flow-control <asym-tx | off | sym | sym-asym-rx>	Prevents one port from sending data faster than the receiving port can handle it. When the receiving port has all the data it can handle, it sends a “pause” frame to the sender. The sender stops sending data until the pause frame expires. Received (asym-rx), transmitted only (asym-tx), or Off modes are supported; the default mode is off.
auto-neg <on | off>	Determines whether ports negotiate with their link partner to operate with parameters common to both links. This method of auto negotiation follows the IEEE 802.3z standard and provides a way to automatically connect multiple types of devices. By default, auto negotiation is enabled.
Flow Control Advertised (advertised-flow-control)	Determines whether flow control setting is advertised. Default is off.
flow-ctrl <asym-tx | off | sym>	Sets flow control.
duplex <half | full>	Half or full. When the port is set to full duplex, it can transmit and receive data simultaneously. With half duplex the port can transmit or receive data, but not both simultaneously. Default duplex is set to full.
egress-frame-cos-policy <ignore | rcos-to-l2-outer-pcp-map>	Sets the egress frame CoS policy.
egress-untag-vlan <VLAN>	Sets the VLAN for egressing untagged frames.
ingress-fixed-dot1dpri <NUMBER: 0-7>	Sets ingress fixed 802.1D value.

Table 5-1
Administrative and operational attributes for ports

Attribute	Description
Inter-packet gap-size <8 | 12>	Sets the inter-packet gap size. This attribute sets the IPG to 12 or 8 bytes. Configuring a reduced IPG of 8 bytes on a port is not compatible with configuring the same port as a Benchmark port-under-test at the same time. You cannot configure a port, previously configured as a Benchmark port-under-test, to function with an IPG of 8 bytes. Similarly, you cannot configure a port, previously configured to function with an IPG of 8 bytes, as a Benchmark port-under-test.
description <String[128]	Configurable 128 character description of the port. By default, the description is blank.
speed <ten | hundred | gigabit | ten-gig | auto>	Sets the speed.
advertised-speed <ten | hundred | gigabit | tengigabit | ten-hundred-gigabit | hundred-gigabit-tengigabit>	Physical port speed, such as 1 Gbps or 10 Gbps. Not applicable for LAG ports. The default value is auto, which matches the speed to the transceiver speed. Any configured auto negotiation settings are ignored for transceivers that do not support auto negotiation, that is, 100M- and 10G-based transceivers. ***Note:*** If you set a value for the speed attribute, the port stays in that speed and a transceiver mismatch error is displayed if there is a mismatch. This functionality is only supported on Ciena-supported transceivers.
advertised-duplex <half | full | half-full>	Sets the advertised value for duplex.
max-frame-size <NUMBER: 512-12288>	Maximum frame size in bytes allowed to ingress/egress the port. The default value is 1526. Jumbo frames are supported with configurable range from 512-12288. Maximum frame size is also referred to as Maximum Transmission Unit (MTU) size. ***Note 1:*** MPLS traffic does not obey the port MTU on the egress side for all platforms, with the exception of the 3926m, 3928, 3942, 5160 and the 5142. ***Note 2:*** Benchmark reflectors on a 3942 and 5142 do not reflect frames greater than 10,239 bytes. ***Note 3:*** The packet capture feature supports a maximum frame size of 1526 bytes.
Port traffic mirroring	
mirror-encap <none | vlan-tag>	Sets port mirroring encapsulation.
mirror-encap-tpid <8100 | 9100 | 88A8>	Sets port mirroring encapsulation TPID.
mirror-encap-vid <NUMBER: 1-4094>	Sets port mirroring encapsulation VID.

Table 5-1
Administrative and operational attributes for ports

Attribute	Description
mirror-port <on | off>	Turns port mirroring on or off. Default is off.
ingress-mirror <Port>	Sets the mirror port for ingress traffic.
egress-mirror <Port>	Sets the mirror port for egress traffic.
Optic transceiver	
mode <default | rj45 | sfp>	Shows the port connector mode, Copper RJ45, Small Form Factor Pluggable (SFP), or SFP+.
Phy loopback	
loopback <on | off>	Indicates whether internal physical loopback is enabled. By default, internal loopback is off. Internal loopback is supported where data that is destined to egress the port in internal loopback mode is looped back through the switch fabric and out the port on which it came in. The loopback occurs in the PHY. Internal loopback can be enabled on any physical port independently, regardless of VLAN membership. For more information, refer to "Port loopback" on page 5-8. Setting the internal loopback attribute automatically sets the port's learn limit to 0, with an action of 'forward'. When the internal loopback setting is 'off', the configured learn limit is then re-applied. If more than two ports within the same VLAN are configured to participate in an internal loopback test, there is a danger of creating a broadcast storm.
power-over-ethernet <on | off>	Sets power over Ethernet state.
Class of Service (CoS)	
fixed-rcos <NUMBER: 0-7>	Sets the fixed resolved CoS value.
fixed-rcolor <green | yellow>	Sets the fixed resolved color, which is green or yellow.
resolved-cos-map <RCOS to RCOS Map>	Sets the R-CoS to F-CoS map.
resolved-cos-policy <dot1d-tag1-cos | fixed-cos | l3-=dscp-cos>	Sets Resolved CoS Policy.
ingress-to-egress-qmap <RCOS Queue Map>	Sets the R-CoS queue map to use in mapping internal CoS (R-CoS) at the ingress port to a CoS queue at an egress port.

Table 5-1
Administrative and operational attributes for ports

Attribute	Description
frame-cos-map <RCOS to FCOS Map>	Sets the egress RCOS -> FCOS Map to use.
ingress-cos-policy <fixed | leave | ip-prec-inherit | phbg-inherit>	Sets the ingress CoS policy.
resolved-cos-remark-l2 <true | false>	Enables or disables frame layer 2 remarking.
VLAN specific	
egress-untag-vlan <VLAN>	Sets the VLAN for egressing untagged data frames.
pvid <VLAN>	Port VLAN ID. Default is 1.
Ingress VLAN Filter (vlan-ingress-filter)	Filters frames that are not members of a configured VLAN. Default is enabled.
Virtual switch and virtual circuit specific	
untagged ctrl vs <Virtual Switch Name>	Untagged control frame virtual switch.
Untagged Data VS (untagged-data-vs)	Untagged data frame virtual switch.
untagged-data-vid <VLAN>	Pushes and pops the specified VID as a Customer VID for frames forwarded to a virtual switch. Applicable only to ports associated with virtual switches.
vs-ingress-filter <on | off>	Determines whether frames that do not explicitly match one of these criteria are assigned to a virtual switch: • subscriber VLAN • untagged-data • untagged-ctrl-vs
vs-egress-filter <on | off>	Works with the enhanced vs-l2-transform mode to filter traffic from NNI to UNI matching SVLAN only

Table 5-1
Administrative and operational attributes for ports

Attribute	Description
vs-ingress-filter-strict <on | off>	Sets strict ingress filter enforcement
vlan-ingress-filter <on | off>	Sets VLAN ingress filter.
vs-l2-transform <i-push, e-pop | i-push, e-pop:stamp | i-stamp:push, e-match-pop:stamp>	Enables VLAN translation for Q-in-Q with virtual switch L2 transform actions.

Small maximum frame sizes

Protocols may fail if the maximum frame size configured is too small. Lowering the maximum frame size must be done cautiously and with the knowledge of the traffic intended on the port to prevent loss of packets. This table provides guidance for the suitability of small maximum frame sizes for common protocols.

Table 5-2
Maximum frame size and protocols

Protocol	512 bytes maximum frame size	1526 bytes (default maximum frame size
LACP	Yes	Yes
LLDP	Yes	Yes
ICMP (IP interface ping)	Yes	Yes
OSPF	Yes	Yes
ISIS	No	Yes
Remote Interface (Telnet)	No	Yes

Protocol restrictions

ISIS does not work on a port when the maximum frame size is lowered from the default setting.

It is not recommended to lower the maximum frame size on any port intended to remotely manage a switch. This is likely to result in Telnet sessions that appear to be hung when a moderately large amount of data is to be displayed.

Using Telnet to a port with a lower maximum frame setting opens a connection. Certain CLI commands may work, but transferring a large amount of data may cause the Telnet session to hang. If this happens, terminate and restart the Telnet session. The max-frame-size can be reset to a larger value to regain admin access to the node.

Received Low Power Detection

Received Low Power Detection is a software polling mechanism used to monitor a system's pluggable-optical ports for fluctuations in optical power. Port "squelching" is the ability for the system to operationally disable any ports that report a low-power alarm condition which signifies that the seated optic is moving towards a loss of signal (LOS). Polling takes place at 5-second intervals for every pluggable port, but to alleviate squelch latency, unpopulated SFP cages are skipped allowing populated SFP cages to be serviced at more frequent intervals. If software polling determines that the received optical power on a port has attenuated below the current low-power alarm threshold for the seated optic, a squelch event is generated. This event identifies the affected port and the reason for the operational state change. A squelch event can also send an SNMP trap containing sufficient information regarding the affected port. A syslog message can also be sent if the device is configured to do so.

An affected port is marked as operationally disabled until the system determines that the seated optic no longer exceeds the low-power alarm threshold in conjunction with at least one of these events:

- system reboot
- re-seat or exchange of the seated SFP
- administratively disabling, followed by enabling of the port

In addition to the manual-intervention mechanisms for bringing a squelched port back to in-service as mentioned previously, a port that has been operationally disabled in this manner can also become operationally enabled without any manual intervention. This occurs when the system determines that the received optical power on the previously 'squelched' port has improved to the point that it has risen above the low-power alarm threshold by a margin of 2dB.

While a port is operationally disabled, the port's operational status reflects that it is down because of a low-power condition by reporting a "squelch" state. In this state, a peer device at the other end of the link does not have a link up indication. A port continues to be serviced by the polling mechanism regardless of whether it is in a normal operational state or marked as operationally down because of a low-power condition.

The procedure for Received Low Power Detection is:

- "Enabling and disabling Received Low Power Detection" on page 5-26

Port loopback

The Internal loopback figure shows internal loopback on Port1. Traffic ingresses Port2 and is intended to egress Port1. The traffic ingresses the switch fabric, is learned, ingresses Port1 and is looped back and sent into the switch fabric.

Figure 5-1
Internal loopback

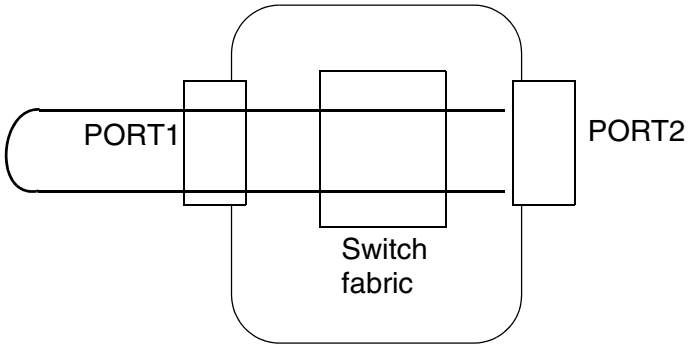

Note: Port loopback is supported on UNI ports only. It is not supported on NNI ports.

Port hold-off

Port hold-off provides a mechanism for dampening the port status by not advertising physical link transitions until the hold-off state duration has elapsed.

A physical link that goes from up to down triggers the hold-off state. Every physical link transition that occurs during the hold-off period is ignored. Once the hold-off timer expires, if the physical link state is still down, then the SAOS begins to advertise the port as down.

During the time period when the port is in hold-off state, the port's operational status in SAOS remains up and its "link state duration" also does not change.

The hold-off state and hold-off timer (in deciseconds or ds), can be configured on physical ports using CLI and SNMP.

No SNMP notification or show CLI is supported to inform that a port is active in the hold-off state, nor are any statistics supported for port hold-off.

An informational syslog raises whenever a port enters or leaves a hold-off state.

Port hold-off has no impact on the status of the LED of the port during physical link transitions. During physical link transitions, the LED of the port has the same behavior as prior to port hold-off.

Once port hold-off is operational on any port, it returns to its normal operational state when:

- The hold-off timer of the port expires.
- There is manual intervention to disable and enable the administrative state of the port. This triggers the port to terminate the hold-off state without waiting for the timer to expire.
- There is manual intervention to disable the port hold-off state or to modify the hold-off timer. This triggers the port to terminate the hold-off state without waiting for the timer to expire.
- The operator changes the mode of a port or physically removes the transceiver—this manual intervention terminates the port hold-off state and restores the port state to a normal operational state.
- The auto-negotiation setting on the port is changed.

Port hold-off on any port impacts mac-address flush behavior. This implies that during the physical link transition to a down state, the dynamic mac-addresses learned on a port are not flushed during the time the port remains in a hold-off state. Once port hold-off is operational on any port, mac-address flush behavior is as follows:

- If the port physical link state is up once the hold-off timer expires or is canceled, then mac-address entries are not flushed.
- If the port physical link state is down once the hold-off timer expires or is canceled, then mac-address entries are flushed because the SAOS begins to advertise this port as down.

Existing link-up and link-down per port statistics are not incremented for a port in a hold-off state. Once the port hold-off timer expires or is canceled, if the physical link is still down, then link-down statistics are incremented.

Feature interactions and limitations

- Port hold-off is supported on all front panel ethernet combo and non-combo physical ports. No support is provided for TDM ports.
- Port hold-off is supported on all Ciena-supported transceivers on which an Ethernet port connected to a transceiver goes down whenever a physical link is not available. Port hold-off is not applicable to transceivers such as TDM SFP, on which the Ethernet side of a port is always up.
- Port hold-off is mutually exclusive to the link-flap feature.
- Port hold-off is mutually exclusive to a system-reserved port/oam-services-role (5160).

- Port hold-off is not supported on aggregation ports.
- Port hold-off is supported on individual physical ports that are members of any aggregate.
- The mechanism port hold-off provides to soak/damp the physical link transitions only at the SAOS and the actual physical link (PHY) downtime depends on the time-frame a physical link is not available.
- The time taken to bring up any port depends on the PCS synchronization (sync) time of PHY used for the port. The PCS sync time of PHY differs depending on these factors:
  - Port types with different PHYs being used across the SAOS product family
  - Media type of the port (optical/RJ-45)
  - Presence of external PHY
  - Different types of connector (SFP, SFP+, XFP, and so on)

  On a port on which hold-off is enabled, all control and data traffic is impacted until the time the actual physical link is not available (downtime plus PCS sync time) due to any reason, such as L1 error, port down, and so on.
- For all the control protocols (such as LACP, LLDP, MPLS, and so on), the port hold-off delays their state machine transitions, if port hold-off is soaking the physical link transitions. That is, without hold-off enabled on a port, all these protocols instantly transition to a new state as soon as the physical link goes down. If port hold-off is enabled on the port, the physical link transition is delayed, thus impacting all such protocol state machine transitions.

  For example, suppose the LACP protocol (having LACP timeout as 30 sec) is enabled on any port P1. If the port P1 physical link goes down for 1 sec, then:
  - If port hold-off is disabled on port P1, then the LACP state machine is notified of link down instantly and the LACP state machine transitions occur.
  - If port hold-off is enabled on port P1 with a hold-off time of 5 ds, the LACP state machine is notified of link-down after 5 ds, once the hold-off timer expires.
  - If port hold-off is enabled on port P1 with a hold-off time of 20 ds, this results in soaking the complete physical link transition (1 sec or 10 ds) and the LACP state machine will not be notified at all.
- For all the control protocols (such as BFD, G.8032, CFM, and so on) which have hello/keep-alive timers shorter than the hold-off timer configured on the port and the physical link detection time, port hold-off has no impact

on their state machine transitions. Also, any link-level actions (if any) triggered through such control protocols' time-outs are not soaked by port hold-off.

For example, suppose BFD (offloaded in the hardware) with a 10 ms timeout is enabled on an IP interface configured on port P1. If the port P1 physical link goes down for 1 sec, then:

— If port hold-off is disabled on port P1, then on BFD timeout event (after 10 ms), the BFD state machine marks the BFD session as down and informs the IP interface the BFD session has timed out.

— If port hold-off is enabled on port P1 with a hold-off time of 20 ds, then even in this case on BFD timeout event (after 10 ms), the BFD state machine marks the BFD session as down and informs the IP interface the BFD session has timed out.

Procedures for configuring ports are:

- "Setting port attributes" on page 5-19
- "Resetting port attributes to default" on page 5-23
- "Disabling a port" on page 5-24
- "Enabling a port" on page 5-25
- "Displaying port attributes" on page 5-27
- "Displaying blade information" on page 5-43
- "Displaying port capabilities" on page 5-47
- "Displaying port Ethernet configuration" on page 5-49
- "Displaying port status" on page 5-50

Port statistics

This table describes port statistics.

Table 5-3
Port statistics

Port statistic	Description
RxBytes	Number of bytes received including those in bad packets.
RxPkts	Number of packets received including all unicast, multicast, broadcast, MAC control and bad packets.
RxCrcErrorPkts	Number of packets received which contained an FCS error and were between 64 and 1518 bytes in length. This value is 1522 bytes if VLAN tagged.

Table 5-3
Port statistics

Port statistic	Description
RxMcastPkts	Number of good multicast packets received that were between 64 and 1518 bytes in length. Excludes MAC control frames. This value is 1522 bytes if VLAN tagged.
RxBcastPkts	Number of good broadcast packets received that were between 64 and 1518 bytes in length. This value is 1522 bytes if VLAN tagged. Excludes MAC control frames.
RxUcastPkts	Number of good unicast packets received that were between 64 and Max Frame Size bytes in length. Excludes MAC control frames.
UndersizePkts	Number of packets received that were less than 64 bytes long and contained a valid FCS and were otherwise well-formed.
OversizePkts	Number of packets received that were longer than 1518 bytes to Max Frame Size and contained a valid FCS and were otherwise well formed.This value is 1522 bytes if VLAN tagged. Includes unicast, multicast and broadcast packets.
FragmentsPkts	Number of packets received that were between 10 and 63 bytes in length and had either an FCS error or an alignment error.
JabbersPkts	Number of packets received that were longer than 1518 bytes to Max Frame Size and had an FCS error or an alignment error. This value is 1522 bytes if VLAN tagged.
RxPausePkts	Number of received valid pause packets that were between 64 and 1518 bytes in length.
RxDropPkts	The total number of valid packets received which were discarded due to lack of resources, that is, rx buffer hits the discard limit, buffer pool full or back pressure discard. RFC 2819 specifies that this number is not necessarily the number of packets dropped; it is just the number of times this condition has been detected.
RxDiscardPkts	The Count of valid frames received which were discarded (filtered) by the Forwarding Process. This includes packets dropped due to lack of resources (RxDropPkts).
RxLOutRangePkts	Number of packets received which exceeded Max Frame Size in length and contained a valid or invalid FCS.
RxInErrorPkts	Number of packets received which have FCS errors, or are either Undersize or Out of Range.
64OctsPkts	Number of packets received that were 64 bytes in length.
65To127OctsPkts	Number of packets received that were between 65 and 127 bytes in length.
128To255OctsPkts	Number of packets received that were between 128 and 255 bytes in length.

Table 5-3
Port statistics

Port statistic	Description
256To511OctsPkts	Number of packets received that were between 256 and 511 bytes in length.
512To1023OctsPkts	Number of packets received that were between 512 and 1023 bytes in length.
1024To1518OctsPkts	Number of packets received that were between 1024 and 1518 bytes in length.
1519To2047OctsPkts	Number of packets received that were between 1519 and 2047 bytes in length.
2048To4095OctsPkts	Number of packets received that were between 2048 and 4095 bytes in length.
4096To9216OctsPkts	Number of packets received that were between 4096 and 9216 bytes in length.
TxBytes	Number of bytes transmitted including those in bad packets.
TxPkts	Number of packets transmitted including all unicast, multicast, broadcast, MAC control and bad packets.
TxExDeferPkts	Number of transmitted packets which experienced multiple deferrals (2 or more deferrals.
TxDeferPkts10	Number of transmitted packets which were deferred on the first transmission but did not experience any subsequent collisions during transmission.
TxGiantPkts	Number of packets transmitted that were longer than 1518 bytes and were otherwise well formed (valid FCS). This value is 1522 bytes if VLAN tagged.
TxUnderRunPkts	Number of transmitted underrun packets.
TxCrcErrorPkts	Number of transmitted packets which contained an FCS error.
TxLCheckErrorPkts	Number of transmitted length check packets
TxLOutRangePkts	Number of transmitted length out of range packets
TxLateCollPkts	Number of transmitted packets which experienced a late collision more than 512 bit times during a transmission attempt
TxExCollPkts	Number of transmitted packets which experienced 16 collisions during transmission and was aborted.
TxSingleCollPkts	Number of transmitted packets which experienced a single collision.
TxCollPkts	Number of transmitted packets which experienced 2-15 collisions (including any late collisions) during transmission.

Table 5-3
Port statistics

Port statistic	Description
TxPausePkts	Number of valid pause control packets transmitted that were between 64 and 1518 bytes in length. This value is 1522 bytes if VLAN tagged.
TxDiscardPkts	Number of transmitted packets dropped on a port due to any reason.
TxUcastPkts	Number of good unicast packets transmitted that were between 64 and 1518 bytes in length. This value is 1522 bytes if VLAN tagged.
TxMcastPkts	Number of good multicast packets transmitted that were between 64 and 1518 bytes in length. This value is 1522 bytes if VLAN tagged.
TxBcastPkts	Number of good broadcast packets transmitted that were between 64 and 1518 bytes in length. This value is 1522 bytes if VLAN tagged.
Tx64OcPkts	Number of packets transmitted that were 64 bytes in length.
Tx65To127OcPkts	Number of packets transmitted that were between 65 and 127 bytes in length.
Tx128To255OcPkts	Number of packets transmitted that were between 128 and 255 bytes in length.
Tx256To511OcPkts	Number of packets transmitted that were between 256 and 511 bytes in length.
Tx512To1023OcPkts	Number of packets transmitted that were between 512 and 1023 bytes in length.
Tx1024To1518OcPkts	Number of packets transmitted that were between 1024 and 1518 bytes in length.
Tx1519To2047OcPkts	Number of packets transmitted that were between 1519 and 2047 bytes in length.
Tx2048To4095OcPkts	Number of packets transmitted that were between 2048 and 4095 bytes in length.
Tx4096To9216OcPkts	Number of packets transmitted that were between 4096 and 9216 bytes in length.
9217to16383OctsPkts	Number of packets received that were between 9217 and 16383 bytes in length.
Tx9217To16383OcPkts	Number of packets transmitted that were between 9217 and 16383 bytes in length.

Procedures for statistics are:

- “Displaying port statistics” on page 5-32
- “Monitoring port statistics” on page 5-37

- “Clearing current statistics” on page 5-42

Transceivers

This section describes

- “Identification”
- “Diagnostics”

Identification

Ciena devices support transceivers that contain a standard serial erasable programmable read-only memory (EPROM) that provides information on the type of SFP used. This information is read from the EPROM:

- Identifier Type, for example, GBIC, SFP
- Extended Identifier Type
- Connector Type, for example, SC, LC, MU, SG
- Vendor Name
- Vendor Organizational Unique Identifier (OUI)
- Vendor Part Number
- Vendor Serial Number
- Vendor Revision Number
- Encoding Algorithm, for example, NRZ, Manchester
- Manufacturing Date Code
- Transceiver Code
- Transceiver SFF-8472 Compliance Version

Diagnostics

Ciena devices support advanced transceivers that have an additional diagnostic serial EPROM. The system software determines if the transceiver has the diagnostic EPROM and provides this information to the user:

- Wavelength/Frequency
- Temperature
- Rx Power
- Tx Power
- Tx Disable State
- Tx Fault State
- Rx Rate Select State

The information is stored in a table by port. The standard EPROM information is updated during initialization or when a new transceiver has been inserted. The diagnostic information is updated at a rate of one port every five seconds. However, this process has a low priority, and in times of a heavy CPU load, the information may be refreshed slowly or not at all.

Transceivers that support diagnostics can trigger events and SNMP traps.

Diagnostics on SFPs are:

- BiasHigh
- BiasLow
- RxPowerHigh
- RxPowerLow
- TempHigh
- TempLow
- TxPowerHigh
- TxPowerLow
- VccHigh
- VccLow

Diagnostics on XFPs are:

- BiasHigh
- BiasLow
- RxPowerHigh
- RxPowerLow
- TempHigh
- TempLow
- TxPowerHigh
- TxPowerLow

Each of these events and traps include a warning and alarm version of each, for example, BiasHighAlarm and BiasHighWarning. Thresholds are set by the SFP vendors, and are not programmable. Both event classes, that is, alarm and warning, are logged under the xcvr-mgr. The warnings are logged using the debug category and warning severity. The alarms are logged using the debug category and minor severity.

These alarms and warnings are based on flags that are set or cleared inside the SFP. These flags are polled at a low priority and slow rate, so flags may be set and then cleared without generating a trap or event. For example, if the TempHigh alarm threshold is exceeded for a few seconds, and then cleared before the flag is polled, it does not trigger a TempHigh alarm. You can forcibly clear alarms and warnings by removing and then reinserting the transceiver or disabling and then enabling the port.

This table lists transceiver states and provides a description of each state.

Table 5-4
Transceiver states

State	Description
INV!	Invalid. The transceiver port state cannot be determined due to a system error.
UCTF	Uncertified. The transceiver is not in the officially supported set of transceivers on that device, for that software version. The transceiver may or may not function properly.
WARN	There are one or more warnings associated with the transceiver, such as the port configuration settings (e.g. speed, autonegotiation, etc) do not match the capabilities of the transceiver. For example, the port is configured for Gig speed, but the transceiver is 100m. Information about which settings are incompatible are available in the output of the port show port command.
FLT!	Fault. The transceiver has been faulted for some reason; typically this is due to EEPROM checksum and/or read failures.
Ena	Enabled
Dis	Disabled

Forward Error Correction mode configuration support

Tunable OTN/FEC SFP+ transceivers (XCVR-TFEC01) provide forward error correction (FEC) capability and configurable transceiver FEC mode (GFEC or EFEC) support.

In SAOS 6.17.1, support for XCVR-TFEC01 was provided with the FEC state as enabled and the FEC mode as GFEC by default. No configuration support was provided for both FEC state and FEC mode.

XCVR-TFEC01 supports two different OTN FEC coding options:

- Generic Forward Error Correction (GFEC) is based on a Reed-Solomon (255,239) code, and it provides up to 6.2 dB improvement in Signal-To-Noise Ratio (SNR).
- Enhanced Forward Error Correction (EFEC) is based on a Reed-Solomon outer code and a Bose Chaudhuri Hocquenghem (BCH) inner code, and it provides a coding gain of up to 8 dB.

In SAOS 6.18, support for configuring the transceiver's FEC mode (GFEC or EFEC mode) has been added. No support for configuring the FEC state has been provided because XCVR-TFEC01 firmware doesn't support this.

A configurable FEC mode attribute has been added to the port XCVR CLI and it includes the following points:

- The FEC mode configured in the software using CLI/SNMP is applied immediately to the transceiver and no SAOS reboot is required for the configuration. Also, the configuration of the FEC mode saved by the user is applied after the system reboot.
- The FEC mode of a transceiver can be configured to either GFEC or EFEC mode and no CLI error is thrown while configuring the FEC mode if there is no transceiver plugged in the port with a 10G capability.
- If the port contains a transceiver that does not support FEC mode capability, then a CLI error is thrown while trying to configure the FEC mode of the transceiver.
- When a transceiver (XCVR-TFEC01) that supports FEC mode capability is inserted, then the configured administrative FEC mode in SAOS/ software is applied to the transceiver.
- If the FEC mode is modified on a link that is operational, then the transceiver (XCVR-TFEC01) gets reset, which results in toggling of the link.

Procedures for configuring transceivers are:

- "Displaying a list of supported optics" on page 5-51
- "Displaying transceiver information" on page 5-53
- "Determining transceiver speed" on page 5-57
- "Tuning XFP transceivers" on page 5-59
- "Tuning OTN FEC SFP+ transceivers" on page 5-61
- "Setting the port connector mode" on page 5-64

Procedure 5-1
Setting port attributes

Set port attributes. For information about port attributes, see “Administrative and operational attributes for ports”.

Step	Action
1	Set port attributes: `port set port <port> {[acceptable-frame-type <all|tagged-only|untagged-only>], [advertised-flow-control <asym-tx|off|sym-asym-rx>],[auto-neg <on|off>], [duplex <half|full>], [description <String[31]>], [egress-frame-cos-policy <ignore|rcos-to-l2-outer-pcp-map>] [egress-untag-vlan <Vlan>], [egress-mirror <PortName>], [fixed-rcos <NUMBER: 0-7], [fixed-rcolor <green|yellow>], [flow-ctrl <asym-rx|asym-tx|off|sym>], [hold-off <on|off>], [hold-off-time <NUMBER: 3-20>], [inter-packet-gap-size <12|8>], [ingress-mirror <PortName>], [link-flap-detect <on|off>], [link-flap-count <NUMBER: 1-64>], [link-flap-detect-time <NUMBER: 1-600>], [link-flap-hold-time <NUMBER: 0-600>], [ingress-to-egress-qmap <RcosQueueMap>] [max-frame-size <NUMBER: 512-12288>], [mirror-port <on|off>], [mode <default|rj45|sfp>], [loopback <on|off>],[pvid <Vlan>], [resolved-cos-policy <dot1d-tag1-cos|fixed-cos|l3-dscp-cos>], [resolved-cos-map <FcosRcosMap>], [resolved-cos-remark-l2 <true|false>], [speed <ten|hundred|gigabit|auto>], [untagged-ctrl-vs <VirtualSwitchName>], [untagged-data-vs <VirtualSwitchName>], [untagged-data-vid <VlanId>], [vlan-ingress-filter <on>], [vs-ingress-filter <on>], [vs-l2-transform <i-push,e-pop|i-push,e-pop:stamp|i-stamp:push,e-match-pop:stamp>][power-over-ethernet <on|off>]}``

where	
port <port>	is the port(s) to set.
[acceptable-frame-type <all|tagged-only|untagged-only>]	sets the acceptable frame types mode as follows: • all - all frames are allowed to ingress, regardless of their tag status. • tagged-only - only tagged frames are allowed to ingress. • untagged-only - only untagged frames are allowed to ingress. ***Note:*** When acceptable frame types are configured on a non-virtual switch UNI, untagged0only mode filters out all three common TPIDs of 0x8100, 0x8898, and 0x9100. For a virtual switch UNI, only the TPID of 0x8100 is recognized as a tagged frame.
[advertised-flow-control <asym-tx|off|sym-asym-rx>	sets the flow.
[auto-neg <on|off>]	sets auto-negotiation.
[duplex <half|full>]	sets duplex.
[description <String[31]>]	sets the port description,
[egress-frame-cos-policy <ignore|rcos-to-l2-outer-pcp-map>]	sets egress frame CoS policy.
[egress-untag-vlan <Vlan>]	sets VLAN for egressing untagged frames.
[egress-mirror <PortName>]	sets egress port mirroring.
[fixed-rcos <NUMBER: 0-7]	sets the fixed resolved CoS.
[fixed-rcolor <green|yellow>]	sets the fixed resolved color.
[flow-ctrl <asym-rx|asym-tx|off|sym>]	sets flow control.
[hold-off <on|off>]	sets the port hold-off state. Default is off.

where	
[hold-off-time <NUMBER: 3-20>]	sets the port hold-off time in deciseconds (1 ds = one-tenth of a second or 100 ms). Default is 10 ds.
[inter-packet-gap-size <12|8>]	sets the inter-packet gap (IPG) size. This attribute sets the IPG to 12 or 8. Configuring a reduced IPG of 8 bytes on a port is not compatible with configuring the same port as a Benchmark port-under-test at the same time. You cannot configure a port, previously configured as a Benchmark port-under-test, to function with an IPG of 8 bytes. Similarly, you cannot configure a port, previously configured to function with an IPG of 8 bytes, as a Benchmark port-under-test/
[ingress-mirror <PortName>]	sets ingress port mirroring,
[link-flap-detect <on|off>]	sets the link flap detection state. Default is off.
[link-flap-count <NUMBER: 1-64>]	sets the link transition count that causes a link flap event.
[link-flap-detect-time <NUMBER 1-600>]	sets the link flap detection window in seconds. The default is 10.
[link-flap-hold-time <NUMBER 0-600>>]	sets the link flap port ren-enable time in seconds. The default is 300.
[ingress-to-egress-qmap <RcosQueueMap>]	sets the ingress to egress queue map.
[max-frame-size <NUMBER: 512-12288>]	sets the maximum frame size. ***Note:*** Benchmark reflectors on 3942 and 5142 devices do not reflect frames greater than 10,239 bytes.
[mirror-port <on|off>]	sets port mirroring state.
[mode <default|rj45|sfp>]	sets the physical interface connector mode.
[loopback <on|off>]	sets the physical loopback state.
[pvid <Vlan>]	sets the PVID.
[resolved-cos-policy <dot1d-tag1-cos|fixed-cos|l3-dscp-cos>]	sets the resolved CoS policy.

where	
[resolved-cos-map <FcosRcosMap>]	sets FCOS -> RCOS map to use.
[resolved-cos-remark-l2 <true|false>]	sets L2 remark based on L3.
[speed <ten|hundred|gigabit|auto>]	sets the speed.
[untagged-ctrl-vs <VirtualSwitchName>]	sets the virtual switch for untagged control frames.
[untagged-data-vs <VirtualSwitchName>]	sets virtual switch for untagged data frames.
[untagged-data-vid <VlanId>]	is the push/pop specified VLAN ID for untagged data frames.
[vlan-ingress-filter <on>]	sets the VLAN ingress filter.
[vs-ingress-filter <on>]	sets the virtual switch ingress filter.
[vs-l2-transform <i-push,e-pop|i-push,e-pop:stamp|i-stamp:push,e-match-pop:stamp>]}	sets the L2 transformation for virtual switch subscribers.
[power-over-ethernet <on|off>	controls the Power-over-Ethernet (PoE) hardware module state. Default is off.

—end—

Example

This example sets these port attributes on port 2.1:

- disables auto-negotiation
- sets the description to 1234_West_Street

```
port set port 2.1 auto-neg off description
1234_West_Street
```

Procedure 5-2
Resetting port attributes to default

Reset port attributes to default values.

Step	Action

1 Reset port attributes to default values:

```
port unset port <port> {hold-off} {hold-off-time} {inter-packet-gap-size} {description} {egress-mirror} {ingress-mirror} {mirror-encap} {mirror-encap-vid} {mirror-encap-tpid} {untagged-ctrl-vs} {untagged-data-vs} {untagged-data-vid} {advertised-speed} {advertised-duplex}
```

where

port <port>	is the port.
hold-off	is the state of the port.
hold-off-time	is the length of the hold-off period.
inter-packet-gap-size	unsets the size of the inter-packet gap.
description	is the port description.
egress-mirror	is egress port mirroring.
ingress-mirror	is ingress port mirroring.
mirror-encap	is port mirroring encapsulation.
mirror-encap-vid	is the port mirroring encapsulation VID.
mirror-encap-tpid	is the port mirroring encapsulation TPID.
untagged-ctrl-vs	is the virtual switch for untagged control frames.
untagged-data-vs	is the virtual switch for untagged data frames.
untagged-data-vid	is the push/pop of VLAN ID for untagged data frames.
{advertised-speed}	is the advertised value for speed.
{advertised-duplex}	is the advertised value for duplex.

—end—

Example

This example clears the description associated with port 1.

```
port unset port 1 description
```

Procedure 5-3
Disabling a port

When you disable a port, the Link State administrative status is changed to disabled and the operational status shows disabled when the link is down.

When you disable a port directly, the transceiver is disabled.

Note: When the port supports PoE, the disable command does not disable the PoE module corresponding to the port-object.

Step	Action
1	Enter this command: `port disable port <PortNameList>` where port <PortNameList> is the port to be disabled.

—end—

Example

This example disables port 1.

```
port disable port 1
```

Procedure 5-4
Enabling a port

When you enable a port, the Link State administrative status is changed to enabled and the operational status shows enabled when the link is up.

When you enable a port directly, the transceiver is enabled.

Step	Action
1	Enter this command: `port enable port <port>` where port <port> is the port to be enabled.

—end—

Example

This example enables port 1.

```
port enable port 1
```

Procedure 5-5
Enabling and disabling Received Low Power Detection

You can

- enable Received Low Power Detection
- disable Received Low Power Detection

Step	Action

To enable Received Low Power Detection

1 Enable Received Low Power Detection:

```
port rx-low-power-detect enable
```

To disable Received Low Power Detection

2 Disable Received Low Power Detection:

```
port rx-low-power-detect disable
```

—end—

Procedure 5-6
Displaying port attributes

Display port attributes to

- verify configuration
- check link status
- troubleshoot issues related to the port

The “port show” commands for specific ports indicate the operational state of the `rx-low-power-detect` feature.

- LOS — loss of signal
- Squelched — operationally disabled
- Blank — the feature is either enabled (and monitoring) or disabled

Step	Action

1 Display port attributes by entering this command:

```
port show [port <port>]
```

where

port <port> is a 32-character string representing the name of the physical port or LAG.

—end—

Example

This example shows sample outputs of the port show command. This example has received low power detection enabled.

```
> port show
+-------------  PORT GLOBAL CONFIGURATION  ---------------+
|           Parameter              |       Value           |
+----------------------------------+-----------------------+
| Rx Low Power Detect Admin State  |  Enabled              |
+----------------------------------+-----------------------+
+-------------------------------------------------------------------------------+
| Port Table        |           Operational Status            |  Admin Config    |
|---------+---------+----+--------------+----+---+-------+----+----+-------+----|
| Port    | Port    |    | Link State   |    |   |       |Auto|    |       |Auto|
| Name    | Type    |Link|   Duration   |XCVR|STP| Mode  |Neg |Link| Mode  |Neg |
|---------+---------+----+--------------+----+---+-------+----+----+-------+----|
| 1       | Gig     |Down|   2d 1h24m26s|    |Dis|       |    |Ena |1000/FD| On |
| 2       | Gig     |Down|   2d 2h10m25s|    |Dis|       |    |Ena |Auto/FD| On |
+---------+---------+----+--------------+----+---+-------+----+----+-------+----+
```

This example has received low power detection disabled.

```
> port show
+-------------  PORT GLOBAL CONFIGURATION  ---------------+
|            Parameter                |       Value       |
+-------------------------------------+-------------------+
| Rx Low Power Detect Admin State |  Disabled               |
+-------------------------------------+-------------------+
+-------------------------------------------------------------------------------+
| Port Table        |           Operational Status            |  Admin Config    |
|-------+---------+----+--------------+----+---+-------+----+----+-------+----|
| Port  | Port    |    |  Link State  |    |   |       |Auto|    |       |Auto|
| Name  | Type    |Link|   Duration   |XCVR|STP| Mode  |Neg |Link| Mode  |Neg |
|-------+---------+----+--------------+----+---+-------+----+----+-------+----|
| 1     | Gig     |Down|   2d 1h24m26s|    |Dis|       |    |Ena |1000/FD| On |
| 2     | Gig     |Down|   2d 2h10m25s|    |Dis|       |    |Ena |Auto/FD| On |
+-------+---------+----+--------------+----+---+-------+----+----+-------+----+
```

This example shows sample output of the port show command for switches that supported system reserved ports.

```
5160> port show
+---------------- PORT GLOBAL CONFIGURATION -----------------+
|             Parameter               |       Value         |
+-------------------------------------+---------------------+
| Rx Low Power Detect Admin State  | Disabled               |
+-------------------------------------+---------------------+
+-------------------------------------------------------------------------------+
| Port Table        |           Operational Status            |  Admin Config    |
|-------+---------+----+--------------+----+---+-------+----+----+-------+----|
| Port  | Port    |    |  Link State  |    |   |       |Auto|    |       |Auto|
| Name  | Type    |Link|   Duration   |XCVR|STP| Mode  |Neg |Link| Mode  |Neg |
|-------+---------+----+--------------+----+---+-------+----+----+-------+----|
| 1     | Gig     |Down|    0d 0h 0m 0s|    |Dis|       |    |Ena |Auto/FD| On |
| 2     | Gig     |Down|    0d 0h 0m 0s|    |Dis|       |    |Ena |Auto/FD| On |
| 17    | G/10Gig |Down|    0d 0h 0m 0s|    |Dis|       |    |Ena |Auto/FD| On |
| 18    | G/10Gig |Down|    0d 0h 0m 0s|    |Dis|       |    |Ena |Auto/FD| On |
| 19    | G/10Gig |Rsvd|    0d 0h 0m 0s|    |Dis|       |    |Ena |Auto/FD| On |
| 20    | G/10Gig |Down|    0d 0h 0m 0s|    |Dis|       |    |Ena |Auto/FD| On |
| 21    | G/10Gig |Down|    0d 0h 0m 0s|    |Dis|       |    |Ena |Auto/FD| On |
| 22    | G/10Gig |Down|    0d 0h 0m 0s|    |Dis|       |    |Ena |Auto/FD| On |
| 23    | G/10Gig |Down|    0d 0h 0m 0s|    |Dis|       |    |Ena |Auto/FD| On |
| 24    | 10Gig   |Down|    0d 0h 0m 0s|    |Dis|       |    |Ena |Auto/FD| On |
+-------+---------+----+--------------+----+---+-------+----+----+-------+----+
```

This example shows sample output of the port show command for a specific reserved port enabled.

```
5160*> po sho po 19
+------------------------------ PORT 19 INFO -------------------------------+
| Field                    | Admin                    | Oper                 |
+--------------------------+--------------------------+----------------------+
| Type                     | Gig/10Gig                | Lag (Gig/10Gig)      |
| L2CFT Status             | Disabled                 | Disabled             |
| L2CFT Profile            |                          |                      |
| Link State               | Enabled                  | Reserved             |
| VLLI State               | Enabled                  |                      |
| Rx Low Power Detection   |                          |                      |
| Ingress ACL              |                          |                      |
| System-Rsvd-Port Status  | Enabled: (vc-transform)  |                      |
+--------------------------+--------------------------+----------------------+
| Aggregation Membership   |                          | ERR                  |
```

```
| VLAN Membership           | 1,127                    |                          |
+---------------------------+--------------------------+--------------------------+
```

This example shows sample output of the port show command for a specific reserved port disabled.

```
5160*> po sho po 19
+------------------------------- PORT 19 INFO --------------------------------+
| Field                     | Admin                    | Oper                     |
+---------------------------+--------------------------+--------------------------+
| Type                      | Gig/10Gig                | Lag (Gig/10Gig)          |
| L2CFT Status              | Disabled                 | Disabled                 |
| L2CFT Profile             |                          |                          |
| Link State                | Enabled                  | Reserved                 |
| VLLI State                | Enabled                  |                          |
| Rx Low Power Detection    |                          |                          |
| Ingress ACL               |                          |                          |
| System-Rsvd-Port Status   | Disabled                 |                          |
+---------------------------+--------------------------+--------------------------+
| Aggregation Membership    |                          | ERR                      |
| VLAN Membership           | 1,127                    |                          |
+---------------------------+--------------------------+--------------------------+
```

This example shows sample output of the port show command for a specific port. The port's operational state is "LOS".

```
> port show port 1

+------------------------------- PORT 1 INFO ---------------------------------+
| Field                     | Admin                    | Oper                     |
+---------------------------+--------------------------+--------------------------+
| Type                      | GigEthernet              |                          |
| Description               |                          |                          |
| Spanning Tree State       | Disabled                 |                          |
| MAC Address               | 00:23:8a:ef:12:e2        |                          |
| Phy Loopback              | Off                      |                          |
| Link Flap Detection       | Off                      |                          |
| Link Flap Count           | 5                        |                          |
| Link Flap Detect Time     | 10                       |                          |
| Link Flap Hold Time       | 300                      |                          |
| Link State                | Enabled                  | Disabled                 |
| VLLI State                | Enabled                  |                          |
| Rx Low Power Detection    |                          | LOS                      |
| State Group Link State    |                          |                          |
| Benchmark State           | Disabled                 | Disabled                 |
+---------------------------+--------------------------+--------------------------+
| Aggregation Membership    |                          |                          |
| VLAN Membership           | 1,127                    | 1,127                    |
+---------------------------+--------------------------+--------------------------+
```

This example shows sample output of the port show command for a specific port. The port's operational state is "Squelched".

```
> port show port 1

+------------------------------- PORT 1 INFO ---------------------------------+
| Field                     | Admin                    | Oper                     |
+---------------------------+--------------------------+--------------------------+
```

```
| Type                        | GigEthernet                | |
| Description                 |                            |
| Spanning Tree State         | Disabled                   |
| MAC Address                 | 00:23:8a:ef:12:e2          |
| Phy Loopback                | Off                        |
| Link Flap Detection         | Off                        |
| Link Flap Count             | 5                          |
| Link Flap Detect Time       | 10                         |
| Link Flap Hold Time         | 300                        |
| Link State                  | Enabled                    | Disabled                  |
| VLLI State                  | Enabled                    |
| Rx Low Power Detection      |                            | Squelched                 |
| State Group Link State      |                            |
| Benchmark State             | Disabled                   | Disabled                  |
+-----------------------------+----------------------------+---------------------------+
| Aggregation Membership      |                            |                           |
| VLAN Membership             | 1,127                      | 1,127                     |
+-----------------------------+----------------------------+---------------------------+
```

This example shows sample output of the port show command for a specific port. The port's operational state is blank, meaning the feature is either disabled or enabled (and monitoring).

```
> port show port 1

+------------------------------- PORT 1 INFO --------------------------------+
| Field                       | Admin                      | Oper                      |
+-----------------------------+----------------------------+---------------------------+
| Type                        | GigEthernet                |                           |
| Description                 |                            |
| Spanning Tree State         | Disabled                   |
| MAC Address                 | 00:23:8a:ef:12:e2          |
| Phy Loopback                | Off                        |
| Link Flap Detection         | Off                        |
| Link Flap Count             | 5                          |
| Link Flap Detect Time       | 10                         |
| Link Flap Hold Time         | 300                        |
| Link State                  | Enabled                    | Disabled                  |
| VLLI State                  | Enabled                    |
| Rx Low Power Detection      |                            |
| State Group Link State      |                            |
| Benchmark State             | Disabled                   | Disabled                  |
+-----------------------------+----------------------------+---------------------------+
| Aggregation Membership      |                            |                           |
| VLAN Membership             | 1,127                      | 1,127                     |
+-----------------------------+----------------------------+---------------------------+
```

This example shows sample output of the port show command for a specific port. The default port hold-off information is included.

```
5142*> port show port 5

+------------------------------- PORT 5 INFO --------------------------------+
| Field                       | Admin                      | Oper                      |
+-----------------------------+----------------------------+---------------------------+
| Type                        | 10GigEthernet              |                           |
| Description                 |                            |                           |
| Spanning Tree State         | Disabled                   |                           |
| MAC Address                 | 9c:7a:03:34:0f:a6          |                           |
| Phy Loopback                | Off                        |                           |
```

Link Flap Detection	Off	
Link Flap Count	5	
Link Flap Detect Time	10	
Link Flap Hold Time	300	
Hold-Off	Disabled	
Hold-Off Time (ds)	10	
Flow Loop Detect State	Off	
Flow Loop Detection	Off	Off
L2CFT Status	Disabled	Disabled
L2CFT Profile		
Link State	Enabled	Disabled
VLLI State		Enabled
Rx Low Power Detection		
State Group Link State		
Degrade Detection Mode	off	off
Degrade State		None
Mode	sfp	unknown
Speed	Auto	0 Gbps
Duplex	full	
Flow Control	off	
Auto Negotiation	Enabled	
Flow Control Advertised	off	
Speed Advertised		
Duplex Advertised		
PVID	1	1
Untag Ingress Data Vid	0	0
Fixed Resolved CoS	0	0
Fixed Resolved Color	green	green
Ingress VLAN Filter	Enabled	Enabled
Ingress VS Filter	Off	Off
Egress VS Filter	Off	Off
Acceptable Frame Type	VLAN tagged only	VLAN tagged only
Egress Untag VLAN	1	1
Max Frame Size	1526	1526
Inter-Packet Gap Size	12	12
Egress Frame Cos Policy	ignore	ignore
Untagged Data VS		
Untagged Ctrl VS		
VS L2 Transform Mode	i-push,e-pop	
Resolved CoS Policy	dot1d-tag1-cos	dot1d-tag1-cos
Service Port Type	Subscriber	Subscriber
Eth VC EtherType	8100	8100
Eth VC EtherType Policy	all	all
Mirror-port	Off	Off
Mirroring Encapsulation	none	
Mirror Encap VID	1	
Mirror Encap TPID	8100	
Ingress-mirror		
Egress-mirror		
Ingress to Egress QMap	Default-RCOS	Default-RCOS
XCVR caps mismatch	speed	
Ingress FCOS->RCOS Map	DefaultFcosRcos	DefaultFcosRcos
Ingress RCOS Remark L2	False	False
Egress RCOS->FCOS Map	DefaultRcosFcos	DefaultRcosFcos
Benchmark State	Disabled	Disabled
Ingress ACL		
Aggregation Membership		
VLAN Membership	1,127	1,127

Procedure 5-7
Displaying port statistics

Two sets of statistics are stored:

- Current statistics, which are the values since the last statistics clear operation.
- Total statistics, which are the values since the last boot-up

The system also calculates throughput values to show current statistics in terms of rate.

You can display

- current statistics
- total statistics
- current throughput statistics
- statistics for specific ports

Note: The port throughput rate measurement is a very rough approximation that must not be expected to match actual rates.

Step	Action

To display current statistics

1 Display current statistics:

```
port show statistics <statistics> [active] [delay
<NUMBER: 1-86400>] [count <NUMBER: 0-4294967295>] [scale
<tera|giga|mega|kilo|none>]
```

where

statistics <statistics>	is statistics collected since the last statistics clear operation.
active	displays active statistics.
delay<NUMBER: 1-86400	is the length of time in seconds to display the statistics.
count <NUMBER: 0-4294967295>	is the number of times to repeat the display.
scale <tera | giga | mega | kilo | none>	is the units for the displayed values.

To display total statistics

2 Display total statistics:

```
port show total-statistics <total-statistics> [active]
[delay <NUMBER: 1..86400>] [count <NUMBER: 0-4294967295>]
[scale <tera|giga|mega|kilo|none>]
```

where

total-statistics <total-statistics>	is all statistics collected since the last boot-up.
active	displays active statistics.
delay<NUMBER: 1..86400	is the length of time in seconds to display the statistics.
count <NUMBER: 0-4294967295>	is the number of times to repeat the display.
scale <tera | giga | mega | kilo | none>	is the units for the displayed values.

To display current throughput statistics

3 Display current throughput statistics:

```
port show throughput <throughput> [active] [delay
<NUMBER: 1..86400>] [count <NUMBER: 0..4294967295>]
[scale <tera|giga|mega|kilo|none>]
```

where

throughput <throughput>	is the port or ports to show.
active	displays active statistics.
delay<NUMBER: 1..86400	is the length of time in seconds to display the statistics.
count <NUMBER: 0..4294967295>	is the number of times to repeat the display.
scale <tera | giga | mega | kilo | none>	is the units for the displayed values. The default value is mega.

To display statistics for specific ports

4 Display statistics for specific ports:

```
port show port <port> [active] [count <NUMBER: 0-
4294967295>] [capabilities] [delay <NUMBER: 1-86400>]
[statistics] [total-statistics] [throughput] [scale
<tera|giga|mega|kilo|none>] [vlan] [type
<all|basic|errors|tx|rx>]
```

where

<port>	is the port or ports that you want to display port statistics for
active	displays active statistics.
count <NUMBER: 0-4294967295>	is the number of repetitions for throughput.
capabilities	displays port capabilities
delay<NUMBER: 1-86400>	is the time between samples in seconds.
statistics	displays port statistics.
total-statistics	displays total port statistics.
throughput	displays port throughput.
scale <tera | giga | mega | kilo | none>	is the scale used to show statistics.
vlan	displays port VLAN membership.
type <all | basic | errors | tx | rx>	is the type of statistics to show.

—end—

Example

This example shows sample output for all active port statistics summary.

```
> port show statistics active

+---------------------------- PORT STATISTICS SUMMARY --------------------+
|   Port   |          Byte                  |          Pkt                 |
|          | Tx              | Rx           | Tx            | Rx           |
+----------+-----------------+--------------+---------------+--------------+
| 1        | 8326248088      | 0            | 67147162      | 0            |
| 2        | 8326247964      | 0            | 67147161      | 0            |
| 12       | 0               | 28879569152  | 0             | 225621634    |
+----------+-----------------+--------------+---------------+--------------+
```

This example shows sample output for all active port total statistics summary.

```
> port show total-statistics active

+----------------------- PORT TOTAL STATISTICS SUMMARY -------------------
|   Port   |          Byte                  |       Pkt                   |
```

```
|         | Tx              | Rx              | Tx           | Rx           |
+---------+-----------------+-----------------+--------------+--------------+
| 1       | 169713065596    | 1053155263708   | 325536280    | 1372943892   |
| 2       | 21452132866     | 19878912        | 77950273     | 310608       |
| 12      | 885114437038    | 193572376072    | 1147664620   | 490131140    |
+---------+-----------------+-----------------+--------------+--------------+
```

This example shows sample output for all active port throughput statistics.

```
> port show throughput active

+--------------- PORT THROUGHPUT SUMMARY 5 SECOND SAMPLE -----------------+
|   Port  |          Bit Rate (Mbps)       |         Pkt Rate (Mpps)        |
|         | Tx              | Rx           | Tx              | Rx           |
+---------+-----------------+--------------+-----------------+--------------+
| 1       |           0.867 |              |           0.001 |              |
| 2       |           0.867 |              |           0.001 |              |
| 12      |                 |        2.008 |                 |        0.003 |
+---------+-----------------+--------------+-----------------+--------------+
```

This example shows sample output for port 1 active statistics.

```
> port show port 1 statistics active

+--------------- PORT 1 STATISTICS -----+
| Statistic           | Value           |
+---------------------+-----------------+
| RxBytes             | 9659942         |
| RxPkts              | 132603          |
| RxCrcErrorPkts      | 1               |
| RxUcastPkts         | 8499            |
| RxMcastPkts         | 81206           |
| RxBcastPkts         | 42897           |
| 64OctsPkts          | 127628          |
| 65To127OctsPkts     | 2345            |
| 256To511OctsPkts    | 2629            |
| 512To1023OctsPkts   | 1               |
| TxBytes             | 1485924         |
| TxPkts              | 12559           |
| TxUcastPkts         | 8233            |
| TxMcastPkts         | 4321            |
| TxBcastPkts         | 5               |
| Tx64OcPkts          | 4648            |
| Tx65To127OcPkts     | 3293            |
| Tx128To255OcPkts    | 4453            |
| Tx256To511OcPkts    | 51              |
| Tx512To1023OcPkts   | 56              |
| Tx1024To1518OcPkts  | 58              |
+---------------------+-----------------+
```

This example shows sample output for port 1 total active statistics.

```
> port show port 1 total-statistics active

+-------------------- PORT 1 STATISTICS --------------------+
| Statistic           | Total Value      | Value           |
+---------------------+------------------+-----------------+
| RxBytes             | 9677152          | 9677152         |
```

```
| RxPkts                | 132855           | 132855           |
| RxCrcErrorPkts        | 1                | 1                |
| RxUcastPkts           | 8635             | 8635             |
| RxMcastPkts           | 81281            | 81281            |
| RxBcastPkts           | 42938            | 42938            |
| 64OctsPkts            | 127873           | 127873           |
| 65To127OctsPkts       | 2350             | 2350             |
| 256To511OctsPkts      | 2631             | 2631             |
| 512To1023OctsPkts     | 1                | 1                |
| TxBytes               | 1493626          | 1493626          |
| TxPkts                | 12654            | 12654            |
| TxUcastPkts           | 8328             | 8328             |
| TxMcastPkts           | 4321             | 4321             |
| TxBcastPkts           | 5                | 5                |
| Tx64OcPkts            | 4731             | 4731             |
| Tx65To127OcPkts       | 3300             | 3300             |
| Tx128To255OcPkts      | 4456             | 4456             |
| Tx256To511OcPkts      | 52               | 52               |
| Tx512To1023OcPkts     | 57               | 57               |
| Tx1024To1518OcPkts    | 58               | 58               |
+-----------------------+------------------+------------------+
```

This example shows sample output for port 1 active throughput statistics.

```
> port show port 1 throughput active

+---------------------- PORT 1 THROUGHPUT --------------------------------+
| Statistic             | Current Value    Delta Value      Rate Mpps & Mbps |
+-----------------------+----------------+----------------+-----------------+
| Time                  |     1:19:56:19 |   0:00:14:20.0 |                 |
| RxBytes               |          9.702 |          0.072 |           0.000 |
| RxPkts                |          0.133 |          0.001 |           0.000 |
| RxUcastPkts           |          0.009 |          0.000 |           0.000 |
| RxMcastPkts           |          0.081 |          0.000 |           0.000 |
| RxBcastPkts           |          0.043 |          0.000 |           0.000 |
| 64OctsPkts            |          0.128 |          0.001 |           0.000 |
| 65To127OctsPkts       |          0.002 |          0.000 |           0.000 |
| 256To511OctsPkts      |          0.003 |          0.000 |           0.000 |
| TxBytes               |          1.500 |          0.027 |           0.000 |
| TxPkts                |          0.013 |          0.000 |           0.000 |
| TxUcastPkts           |          0.008 |          0.000 |           0.000 |
| Tx64OcPkts            |          0.005 |          0.000 |           0.000 |
| Tx65To127OcPkts       |          0.003 |          0.000 |           0.000 |
| Tx128To255OcPkts      |          0.004 |          0.000 |           0.000 |
| Tx256To511OcPkts      |          0.000 |          0.000 |           0.000 |
| Tx512To1023OcPkts     |          0.000 |          0.000 |           0.000 |
| Tx1024To1518OcPkts    |          0.000 |          0.000 |           0.000 |
+-----------------------+----------------+----------------+-----------------+
```

Procedure 5-8
Monitoring port statistics

Use this procedure to continuously monitor port statistics for all ports or for specific ports. The system displays the statistics and automatically clears the screen before displaying the updated values.

To stop monitoring, press Ctrl+C.

Note: Simultaneously monitoring the throughput in more than one session on a 39xx/51xx might cause incorrect throughput results.

Step	Action

To monitor all ports for current statistics

1 Monitor all ports for current statistics:

```
port monitor statistics <statistics> [active] [delay
<NUMBER: 1-86400>] [count <NUMBER: 0-4294967295>] [scale
<tera|giga|mega|kilo|none>]
```

where

statistics <statistics>	displays all statistics collected from the last statistics clear operation.
active	displays active statistics.
delay<NUMBER: 1-86400	is the length of time in seconds to display the statistics.
count <NUMBER: 0-4294967295>	is the number of times to repeat the display.
scale <tera | giga | mega | kilo | none>	is the units for the displayed values.

To monitor all ports for total statistics

2 Monitor all ports for total statistics:

```
port monitor total-statistics <total-statistics> [active]
[delay <NUMBER: 1-86400>] [count <NUMBER: 0-4294967295>]
[scale <tera|giga|mega|kilo|none>]
```

where

total-statistics <total-statistics>	displays all statistics collected since the last boot-up.
active	displays active statistics.
delay<NUMBER: 1-86400	is the length of time in seconds to display the statistics.
count <NUMBER: 0-4294967295>	is the number of times to repeat the display.
scale <tera | giga | mega | kilo | none>	is the units for the displayed values.

To monitor all ports for throughput statistics

3 Monitor all ports for throughput statistics:

```
port monitor throughput <throughput> [active] [delay
<NUMBER: 1-86400>] [count <NUMBER: 0-4294967295>] [scale
<tera|giga|mega|kilo|none>]
```

where

throughput <throughput>	displays current throughput statistics.
active	displays active statistics.
delay<NUMBER: 1-86400	is the length of time in seconds to display the statistics.
count <NUMBER: 0-4294967295>	is the number of times to repeat the display.
scale <tera | giga | mega | kilo | none>	is the units for the displayed values.

To monitor specific ports

4 Monitor specific ports:

```
port monitor port <port> [active] [delay <NUMBER: 1-
86400>] [scale <tera|giga|mega|kilo|none>] {statistics]
[total-statistics] [throughput] [type
<all|basic|errors|tx|rx>]
```

where

port <port>	is the port or ports that you want to monitor
active	displays active statistics.
delay<NUMBER: 1-86400	is the length of time in seconds to display the statistics.
scale <tera | giga | mega | kilo | none>	is the units for the displayed values.
statistics	displays port statistics.
total-statistics	displays total port statistics.
throughput	displays port throughput.
type <all | basic | errors | tx | rx>	is the type of statistics to show.

—end—

Example

This example shows sample output of monitoring total statistics for all active ports.

```
> port monitor total-statistics active delay 10
<Screen clears>

+---------------------- PORT TOTAL STATISTICS SUMMARY ---------------------+
|   Port   |          Byte                    |          Pkt                 |
|          | Tx              | Rx             | Tx            | Rx           |
+----------+-----------------+----------------+---------------+--------------+
| 1        | 170475797780    | 1053155263708  | 331687346     | 1372943892   |
| 2        | 22214864802     | 19878912       | 84101337      | 310608       |
| 12       | 885114437038    | 196217923208   | 1147664620    | 510799477    |
+----------+-----------------+----------------+---------------+--------------+
```

This example shows sample output of monitoring statistics for all ports.

```
> port monitor statistics active delay 10
<Screen clears>

+-------------------------- PORT STATISTICS SUMMARY ----------------------+
|   Port   |          Byte                    |          Pkt                 |
|          | Tx              | Rx             | Tx            | Rx           |
+----------+-----------------+----------------+---------------+--------------+
| 2        | 36938335        | 67861410       | 224689        | 1023888      |
| 3        | 3522850         | 21965574       | 22329         | 228469       |
+----------+-----------------+----------------+---------------+--------------+
```

This example shows sample output of monitoring throughput statistics for all ports.

```
> port monitor throughput active
Info: This CLI output may take a while to display press CTRL-C to abort

<Screen clears>
+------------- PORT THROUGHPUT SUMMARY   5 SECOND SAMPLE -----------------+
|   Port   |          Bit Rate (Mbps)          |       Pkt Rate (Mpps)         |
|          | Tx              | Rx              | Tx             | Rx           |
+----------+-----------------+-----------------+----------------+--------------+
| 1        |           0.860 |                 |          0.001 |              |
| 2        |           0.860 |                 |          0.001 |              |
| 12       |                 |           2.984 |                |        0.003 |
+----------+-----------------+-----------------+----------------+--------------+
```

This example shows sample output of monitoring total statistics for a specific port.

```
> port monitor port 1 total-statistics active delay 10
<Screen clears>

INFO: Waiting 10 seconds for display. Abort with CTRL-c

+--------------------- PORT 1 STATISTICS --------------------+
| Statistic           | Total Value       | Value             |
+---------------------+-------------------+-------------------+
| RxBytes             | 9844868           | 9844868           |
| RxPkts              | 135233            | 135233            |
| RxCrcErrorPkts      | 1                 | 1                 |
| RxUcastPkts         | 9572              | 9572              |
| RxMcastPkts         | 82222             | 82222             |
| RxBcastPkts         | 43438             | 43438             |
| 64OctsPkts          | 130164            | 130164            |
| 65To127OctsPkts     | 2406              | 2406              |
| 256To511OctsPkts    | 2662              | 2662              |
| 512To1023OctsPkts   | 1                 | 1                 |
| TxBytes             | 1580957           | 1580957           |
| TxPkts              | 13396             | 13396             |
| TxUcastPkts         | 9070              | 9070              |
| TxMcastPkts         | 4321              | 4321              |
| TxBcastPkts         | 5                 | 5                 |
| Tx64OcPkts          | 5195              | 5195              |
| Tx65To127OcPkts     | 3519              | 3519              |
| Tx128To255OcPkts    | 4480              | 4480              |
| Tx256To511OcPkts    | 65                | 65                |
| Tx512To1023OcPkts   | 60                | 60                |
| Tx1024To1518OcPkts  | 77                | 77                |
+---------------------+-------------------+-------------------+
```

This example shows sample output for monitoring statistics for a specific port.

```
> port monitor port 1 statistics active delay 10
<Screen clears>

INFO: Waiting 10 seconds for display. Abort with CTRL-c

+--------------- PORT 1 STATISTICS -----+
| Statistic           | Value           |
+---------------------+-----------------+
| RxBytes             | 9866442         |
| RxPkts              | 135531          |
| RxCrcErrorPkts      | 1               |
| RxUcastPkts         | 9648            |
| RxMcastPkts         | 82368           |
| RxBcastPkts         | 43514           |
| 64OctsPkts          | 130448          |
| 65To127OctsPkts     | 2415            |
| 256To511OctsPkts    | 2667            |
| 512To1023OctsPkts   | 1               |
| TxBytes             | 1586934         |
| TxPkts              | 13454           |
| TxUcastPkts         | 9128            |
| TxMcastPkts         | 4321            |
| TxBcastPkts         | 5               |
| Tx64OcPkts          | 5238            |
| Tx65To127OcPkts     | 3530            |
| Tx128To255OcPkts    | 4482            |
| Tx256To511OcPkts    | 66              |
| Tx512To1023OcPkts   | 60              |
| Tx1024To1518OcPkts  | 78              |
+---------------------+-----------------+
```

This example shows sample output for monitoring throughput statistics for a specific port.

```
> port monitor port 1 throughput active
<Screen clears>

INFO: Waiting 5 seconds for display. Abort with CTRL-c

+---------------------- PORT 1 THROUGHPUT ---------------------------------+
| Statistic         | Current Value  | Delta Value       | Rate Mpps & Mbps |
+-------------------+----------------+-------------------+------------------+
| Time              |     1:20:30:17 |     0:00:00:05.0  |                  |
| RxBytes           |          9.887 |             0.000 |            0.000 |
| RxPkts            |          0.136 |             0.000 |            0.000 |
| RxUcastPkts       |          0.010 |             0.000 |            0.000 |
| RxMcastPkts       |          0.083 |             0.000 |            0.000 |
| RxBcastPkts       |          0.044 |             0.000 |            0.000 |
| 64OctsPkts        |          0.131 |             0.000 |            0.000 |
| TxBytes           |          1.593 |             0.000 |            0.000 |
| TxPkts            |          0.014 |             0.000 |            0.000 |
| TxUcastPkts       |          0.009 |             0.000 |            0.000 |
| Tx65To127OcPkts   |          0.004 |             0.000 |            0.000 |
+-------------------+----------------+-------------------+------------------+
```

Procedure 5-9
Clearing current statistics

Clear current statistics when you no longer want to view them. Clearing current statistics does not clear total statistics.

You can clear current statistics for

- all ports
- specific ports

Note: The port clear command does not clear TDM port statistics. For more information about monitoring TDM statistics, refer to “Performance monitoring,” in *39XX/51XX Service Delivery, Aggregation and Virtualization Switches Fault and Performance Management.*

Step	Action

To clear current statistics for all ports

1 Clear current statistics for all ports:

```
port clear statistics
```

To clear current statistics for specific ports

2 Clear current statistics for specific ports:

```
port clear port <port> statistics
```

where

port <port> is the port or ports that you want to clear statistics for.

—end—

Procedure 5-10
Displaying blade information

Display blade information.

Step	Action
1	Display blade information: `blade show [attributes] [capabilities] [information] [state]` where attributes — displays hardware device identification information. Hardware device identification information is device type, hardware version, serial number, MAC address, manufactured date, and E-PROM (param) version. The blade show attributes output varies depending on the hardware platform. The CLEI Code is only displayed for 3916, 3930, 3931, and 5150 devices. capabilities — displays the capabilities of the board and ports, including blade type, RAM and flash file sizes, port types supported, and enhanced ports. information — displays general blade information, including blade type, number of ports, MAC address, administrative and operational states, and date of last reboot. state — displays administrative and operational states.

—end—

Example

This example shows sample output for a single blade device.

```
> blade show

+---------------------- BLADE SUMMARY ----------------------+
| Slot | Ports | PortBaseMac       | BladeType | OperState   |
|------+-------+-------------------+-----------+-------------|
| 1    | 12    | 00:03:18:55:71:d2 | Single    | Enabled     |
+------+-------+-------------------+-----------+-------------+
```

This example shows all extended blade details in one command on a 5150 switch.

```
> blade show attributes capabilities information state

//The blade attributes:
```

```
+----------------- BLADE DEVICE ID -----------------+
| Parameter                  | Value                |
+----------------------------+----------------------+
| Board Device Type          | 091                  |
| Board Hardware Version     | 1705150830/004       |
| Board Serial Number        | B6054200             |
| Board MAC Address          | 00:03:18:ac:e9:40    |
| Manufactured Date          | 11-11-2010           |
| CLEI Code                  | COMP100BRA           |
| Location of Manufacture    | 1                    |
| Module Part Num            | 1705150900/009       |
| Module Serial Num          | M6145938             |
| Param Version              | 007                  |
+----------------------------+----------------------+

+--------------- MODULE 2 DEVICE ID ----------------+
| Parameter                  | Value                |
+----------------------------+----------------------+
| Board Device Type          | 092                  |
| Board Hardware Version     | 1705100810/006       |
| Board Serial Number        | B6061295             |
| Manufactured Date          | 22-10-2010           |
| CLEI Code                  | COUIA3CPAA           |
| Location of Manufacture    | 1                    |
| Module Part Num            | 1705100900/004       |
| Module Serial Num          | M6154551             |
| Param Version              | 007                  |
+----------------------------+----------------------+

+--------------- MODULE 3 DEVICE ID ----------------+
| Parameter                  | Value                |
+----------------------------+----------------------+
| Board Device Type          | 093                  |
| Board Hardware Version     | 1705101810/003       |
| Board Serial Number        | B6004593             |
| Manufactured Date          | 04-07-2010           |
| CLEI Code                  | COUIA7APAA           |
| Location of Manufacture    | 1                    |
| Module Part Num            | 1705101900/002       |
| Module Serial Num          | M6043072             |
| Param Version              | 007                  |
+----------------------------+----------------------+

//The blade capabilities:

+---------------------------- BOARD CAPABILITIES ----------------------------+
| Parameter          | Value                                                 |
+--------------------+-------------------------------------------------------+
| Capability Class   | 0                                                     |
| Board Type         | 091                                                   |
| Board Name         | 5150                                                  |
| Board Description  | 5150 Service Aggregation Switch                       |
| Blade Type         | Single                                                |
| No. Ports          | 52                                                    |
| Has Dcard          | No                                                    |
| Address Ram Size   | 0x0                                                   |
| Boot Flash Size    | 0x200000                                              |
| Packet Ram Size    | 0x0                                                   |
| Program Ram Size   | 0x2000000                                             |
+--------------------+-------------------------------------------------------+
| No. 10 Gig Ports   | 4                                                     |
```

```
| No. Gig Ports       | 48                                                                |
| No. Fe Ports        | 0                                                                 |
| No. 100Fx Ports     | 0                                                                 |
| No. Eth Ports       | 0                                                                 |
| Total Ports         | 52                                                                |
| Enhanced Ports List | 2.1 2.2                                                           |
+---------------------+-------------------------------------------------------------------+
```

SUMMARY PORT CAPABILITIES

Port	Type	Speed	Duplex	Aneg	Pause	Enh
1.1	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.2	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.3	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.4	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.5	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.6	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.7	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.8	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.9	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.10	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.11	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.12	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.13	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.14	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.15	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.16	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.17	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.18	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.19	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.20	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.21	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.22	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.23	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.24	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.25	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.26	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.27	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.28	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.29	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.30	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.31	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.32	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.33	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.34	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.35	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.36	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.37	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.38	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.39	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.40	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.41	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.42	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.43	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.44	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.45	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.46	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.47	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
1.48	100/G	100Mbps,1Gig,Auto	half,full	on,off	off,sym,a-rx	No
2.1	10GigEthernet	10Gig	full	N/A	off	Yes
2.2	10GigEthernet	10Gig	full	N/A	off	Yes
3.1	10GigEthernet	10Gig	full	N/A	off	No
3.2	10GigEthernet	10Gig	full	N/A	off	No

```
//The blade information:
```

BLADE INFO

Parameter	Value
Slot	1
Blade Type	Single

```
| Number of Ports        | 52                          |
| Port Base MAC Address  | 00:03:18:ac:e9:42           |
| Admin State            | Enabled                     |
| Oper State             | Enabled                     |
| Last Reboot            | Thu Jan  1 00:00:00 1970 |
+------------------------+-----------------------------+

//The blade state:

+----- BLADE STATE ------+
| AdminState | OperState |
+------------+-----------+
```

Procedure 5-11
Displaying port capabilities

You can display:

- port capabilities for the chassis and a summary of port capabilities
- capabilities for a specified port

Step	Action

To display port capabilities for the chassis and a summary of port capabilities

1 Display chassis port capabilities and a summary of all port capabilities:

```
port show capabilities
```

To display capabilities for a specific port

2 Display capabilities for a specific port:

```
port show port <port> capabilities
```

where

port <port> is a 32-character string representing the name of the physical port or LAG.

—end—

Example

This example shows sample output for the port show capabilities command.

```
> port show capabilities

+---------------------+-----------------------------------------------------+
| No. 10 Gig Ports    | 4                                                   |
| No. Gig Ports       | 8                                                   |
| No. Fe Ports        | 0                                                   |
| No. 100Fx Ports     | 0                                                   |
| No. Eth Ports       | 0                                                   |
| Total Ports         | 12                                                  |
+---------------------+-----------------------------------------------------+

+------------------------- SUMMARY PORT CAPABILITIES ---------------------------+
| Port| Type         | Speed                | Duplex   | Aneg  | Pause            |
+-----+--------------+----------------------+----------+-------+------------------+
| 1   | 10/100/G     | 10Mbps,100Mbps,1Gig,Auto| half,full| on,off| off,sym,a-rx |
| 2   | 10/100/G     | 10Mbps,100Mbps,1Gig,Auto| half,full| on,off| off,sym,a-rx |
| 3   | 10/100/G     | 10Mbps,100Mbps,1Gig,Auto| half,full| on,off| off,sym,a-rx |
| 4   | 10/100/G     | 10Mbps,100Mbps,1Gig,Auto| half,full| on,off| off,sym,a-rx |
| 5   | 10/100/G     | 10Mbps,100Mbps,1Gig,Auto| half,full| on,off| off,sym,a-rx |
| 6   | 10/100/G     | 10Mbps,100Mbps,1Gig,Auto| half,full| on,off| off,sym,a-rx |
| 7   | 10/100/G     | 10Mbps,100Mbps,1Gig,Auto| half,full| on,off| off,sym,a-rx |
| 8   | 10/100/G     | 10Mbps,100Mbps,1Gig,Auto| half,full| on,off| off,sym,a-rx |
| 9   | 10GigEthernet| 10Gig                | full     | N/A   | off              |
| 10  | 10GigEthernet| 10Gig                | full     | N/A   | off              |
| 11  | 10GigEthernet| 10Gig                | full     | N/A   | off              |
| 12  | 10GigEthernet| 10Gig                | full     | N/A   | off              |
+-----+--------------+----------------------+----------+-------+------------------+
```

This example shows sample output for the port show capabilities command applied to a specified port.

```
> port show port 1 capabilities

+------------------------ PORT 1 CAPABILITIES ---------------------+
| Field                          | Value                           |
+--------------------------------+---------------------------------+
| Port Number                    | 1                               |
| Port Type                      | 10/100/G                        |
| Port Speed                     | 10Mbps,100Mbps,1Gig,Auto        |
| Port Duplex                    | half,full                       |
| Port Auto Negotiation          | on,off                          |
| Port Pause Advertisement       | off,sym,a-tx,s-a-rx             |
| Port Pause                     | off,sym,a-rx                    |
| Port Feature Capabilities      | Normal                          |
+--------------------------------+---------------------------------+
```

Procedure 5-12
Displaying port Ethernet configuration

You can display port attributes for Ethernet configuration for all or for a specific line module, including name, type, admin status, speed, duplex, flow control, flow control advertised, auto negotiation, and MTU size.

Step	Action
1	Display port attributes for Ethernet configuration: `port show ethernet-config`

—end—

Example

This example shows sample output for the port show ethernet-config command.

```
> port show ethernet-config
+------------------------ PORT ETHERNET CONFIGURATION -------------------------+
|                                                                              |
|                                                                    Mirror    |
|  Port     Port      Admin                            FC    Auto   MTU  Status |
|  Name     Type      Status   Speed    Dplx    FC     Adv   Neg    Size State|Eg|Ig|
+--------+--------+--------+-------+------+-----+-----+-----+------+-----+--+--+
|  1       Gig       Ena      1000      ??    off   off    On     1526   Off   0  0|
|  2       Gig       Ena      1000    Full    off   off    On     1526   Off   0  0|
|  3       Gig       Ena      1000    Full    off   off    On     1526   Off   0  0|
|  4      10/100/G   Ena      1000      ??    off   off    On     1526   Off   0  0|
|  5      10/100/G   Ena      1000      ??    off   off    On     1526   Off   0  0|
|  6      10/100/G   Ena      1000      ??    off   off    On     1526   Off   0  0|
|  7      10/100/G   Ena      1000      ??    off   off    On     1526   Off   0  0|
|  8      10/100/G   Ena      1000      ??    off   off    On     1526   Off   0  0|
|  9       10Gig     Ena        10G     ??    off   off    Off    1526   Off   0  0|
|  10      10Gig     Ena        10G     ??    off   off    Off    1526   Off   0  0|
|  11      10Gig     Ena        10G     ??    off   off    Off    1526   Off   0  0|
|  12      10Gig     Ena        10G     ??    off   off    Off    1526   Off   0  0|
+--------+--------+--------+-------+------+-----+-----+-----+------+-----+--+--+
```

Procedure 5-13
Displaying port status

Port status included the operational information, such as the link state, link state duration, whether transceivers are enabled or disabled, speed, duplex, maximum frame size, and flow control. You can display the status for all ports or ports on a specific line module.

Step	Action
1	Display port status: `port show status`

—end—

Example

This example shows sample output for the port show status command.

```
> port show status
+------------------------ PORT OPERATIONAL STATUS -------------------------+
|  |                            |      |Link State|     |    |Speed/  |MTU  |Flow  |
|##|  Description               |Link  |Duration  |XCVR |STP |Duplex  |Size |Ctrl  |
+--+----------------------------+------+----------+-----+----+--------+-----+------+
|1 |                            |Down  |  1d21h35m|     |Dis |        |1526 |      |
|2 |                            | Up   |  1d22h37m|Ena  |FWD |1000/FD |1526 |off   |
|3 |                            | Up   |  1d21h33m|Ena  |FWD |1000/FD |1526 |off   |
|4 |                            |Down  |  0h 0m 0s|     |Dis |        |1526 |      |
|5 |                            |Down  |  0h 0m 0s|     |Dis |        |1526 |      |
|6 |                            |Down  |  0h 0m 0s|     |Dis |        |1526 |      |
|7 |                            |Down  |  0h 0m 0s|     |Dis |        |1526 |      |
|8 |                            |Down  |  0h 0m 0s|     |Dis |        |1526 |      |
|9 |                            |Down  |  0h 0m 0s|     |Dis |        |1526 |      |
|10|                            |Down  |  0h 0m 0s|     |Dis |        |1526 |      |
|11|                            |Down  |  0h 0m 0s|     |Dis |        |1526 |      |
|12|                            |Down  |  0h 0m 0s|     |Dis |        |1526 |      |
+--+----------------------------+------+----------+-----+----+--------+-----+------+
```

Procedure 5-14
Displaying a list of supported optics

Small Form-factor Pluggable (SFP) and 10 Gigabit SFPs (XFP) devices are hot-swappable compact optical transceivers. Port transceiver information, including status and type, is available via CLI or SNMP.

The system software supports transceivers that are compliant with these documents:

- XFP Xcvr spec SFF INF 8077i Rev 4.5, Tunable Xcvr spec SFF-8477 Rev 1.3 Draft.
- Small Form Factor Pluggable (SFP) Transceiver Multi Source Agreement, September 14, 2000
- Digital Diagnostic Monitoring Interface for Optical transceivers SFF-8473, Draft Revision 9.0, April 4, 2002.

Step	Action
1	Display a list of supported optics: `port xcvr show supported` ***Note:*** The output of "port xcvr show supported" is a generic table showing the speed capabilities of each transceiver. However, the actual operational speed depends upon the capabilities that are supported on the specific platform and port.

—end—

Example

This example shows sample output for the port xcvr show supported command.

```
> port xcvr show supported

+----All Supported Transceivers----+
|                    |             |
|    Part Number     | XCVR Speed  |
+--------------------+-------------+
|        XCVR-010X31 |     100M    |
|        XCVR-040X31 |     100M    |
|        XCVR-040R55 |     100M    |
|        XCVR-040R31 |     100M    |
|        XCVR-010S55 |     100M    |
|        XCVR-010S31 |     100M    |
|        XCVR-010L31 |     100M    |
|        XCVR-040L31 |     100M    |
|        XCVR-100D43 |  100M / 1G  |
|        XCVR-100D45 |  100M / 1G  |
|        XCVR-100D47 |  100M / 1G  |
```

```
...                                  |
+---------------------+--------------+
```

Procedure 5-15
Displaying transceiver information

Diagnostics can be displayed for transceivers that support diagnostics. If the transceiver does not have internal diagnostics capabilities, an error is returned.

You can display

- diagnostic information for a specified port
- a summary of transceiver status
- a summary of transceiver status for a specific port
- vendor EPROM data for a specific port

Step	Action

To display diagnostic information for a specified port

1 Display diagnostic information for a specified port:

```
port show port <port> diagnostics
```

where

port <port> is the port or ports to show.

To display a summary of transceiver status

2 Display a summary of transceiver status:

```
port xcvr show
```

To display a summary of transceiver status for a specific port

3 Display a summary of transceiver status for a specific port:

```
port xcvr show port <port> state
```

where

port <port> is the port or ports to show.

To display vendor EPROM data for a specific port

4 Display vendor EPROM data for a specific port:

```
port show port <port> vendor
```

where

port <port> is the port or ports to show.

—end—

Example

This example shows sample output for a transceiver that supports diagnostics.

```
> port xcvr show port 3 diagnostics

+--------------------- XCVR DIAGNOSTICS - Port 3        -----------------+
|                     |           |          Alarm      |        Warning      |
| Output              | Value     | Threshold           | Flag | Threshold    | Flag |
+---------------------+-----------+---------------------+------+--------------+------+
| Temp     (degC)     |   42.031  | HIGH 105.000        | 0    | HIGH 100.000 | 0    |
|                     |           | LOW  -45.000        | 0    | LOW  -40.000 | 0    |
+---------------------+-----------+---------------------+------+--------------+------+
| Vcc     (volts)     |    3.299  | HIGH   3.630        | 0    | HIGH   3.460 | 0    |
|                     |           | LOW    2.970        | 0    | LOW    3.130 | 0    |
+---------------------+-----------+---------------------+------+--------------+------+
| Bias       (mA)     |    4.112  | HIGH  12.000        | 0    | HIGH  10.000 | 0    |
|                     |           | LOW    1.000        | 0    | LOW    2.000 | 0    |
+---------------------+-----------+---------------------+------+--------------+------+
| Tx Power   (mW)     |    0.257  | HIGH   1.412        | 0    | HIGH   0.707 | 0    |
|                     |           | LOW    0.056        | 0    | LOW    0.112 | 0    |
+---------------------+-----------+---------------------+------+--------------+------+
| Tx Power (dBm)      |   -5.8905 | HIGH  +1.4999       | 0    | HIGH  -1.5003| 0    |
|                     |           | LOW  -12.5026       | 0    | LOW   -9.5001| 0    |
+---------------------+-----------+---------------------+------+--------------+------+
| Rx Power   (mW)     |    0.0574 | HIGH   1.9954       | 0    | HIGH   1.0000| 0    |
|                     |           | LOW    0.0100       | 0    | LOW    0.0200| 0    |
+---------------------+-----------+---------------------+------+--------------+------+
| Rx Power (dBm)      |  -12.4109 | HIGH  +3.0001       | 0    | HIGH  +0.0000| 0    |
|                     |           | LOW  -20.0000       | 0    | LOW  -16.9897| 0    |
+---------------------+-----------+---------------------+------+--------------+------+
```

This example shows sample output for a summary of transceiver status.

```
> port xcvr show

+----+-----+-----+---------Transceiver-Status------------+-----+----------------+----+
|    |Admin| Oper|                                       |Ciena|Ether Medium &  |Diag|
|Port|State|State|      Vendor Name & Part Number        | Rev |Connector Type  |Data|
+----+-----+-----+---------------------------------------+-----+----------------+----+
|1.1 |Ena  |     |CIENA-FBX XCVR-A00G85 Rev10            |D    |1000BASE-SX/LC  |Yes |
|1.2 |Ena  |     |CIENA-FBX XCVR-A00G85 Rev10            |D    |1000BASE-SX/LC  |Yes |
|1.3 |Empty|     |                                       |     |                |    |
|1.4 |Empty|     |                                       |     |                |    |
|1.5 |Empty|     |                                       |     |                |    |
|1.6 |Empty|     |                                       |     |                |    |
|1.7 |Empty|     |                                       |     |                |    |
|1.8 |Empty|     |                                       |     |                |    |
|1.9 |Ena  |     |CIENA-OCP XCVR-S10U27 Rev0000          |A    |10GBASE-LR/LC   |Yes |
|1.10|Ena  |     |CIENA-OCP XCVR-S10U33 Rev0000          |A    |10GBASE-LR/LC   |Yes |
|1.11|Ena  |     |CIENA-OCP XCVR-S40U27 Rev0000          |A    |10GBASE-ER/LC   |Yes |
|1.12|Ena  |     |CIENA-OCP XCVR-S40U33 Rev0000          |A    |10GBASE-ER/LC   |Yes |
|1.13|Empty|     |                                       |     |                |    |
...
|2.1 |Ena  |UCTF |JDSU JXPR01LMB81CE Rev01               |     |10GBASE-LW/LR/LC|Yes |
|2.2 |Ena  |     |WORLDWIDEPACKETS XCVR-010V31 Rev1a      |     |10GBASE-LW/LR/LC|Yes |
|3.1 |Ena  |     |CIENA XCVR-010V31 Rev1A                |C    |10GBASE-LW/LR/LC|Yes |
|3.2 |Ena  |     |WORLDWIDEPACKETS XCVR-000Z85 Rev1a      |     |10GBASE SW/SR/LC|Yes |
+----+-----+-----+---------------------------------------+-----+----------------+----+
```

This example shows sample output for a summary of transceiver status for a specific port.

```
> port xcvr show port 12 state

+----+-----+-----+---------Transceiver-Status--------+----------------+----+
|    |Admin| Oper|                                   |Ether Medium & |Diag|
|Port|State|State|      Vendor Name & Part Number     |Connector Type |Data|
+----+-----+-----+-----------------------------------+----------------+----+
|12  |Ena  |Ena  |CIENA XCVR-010Y31 Rev10            |1000BASE-LX/LC |    |
+----+-----+-----+-----------------------------------+----------------+----+
```

This example shows sample output for vendor EPROM data for a specific port.

```
> port xcvr show port 7 vendor

+------------------------ XCVR VENDOR DATA - Port 7    -------------------------+
| Parameter                | Value              | Decoded String Equivalent    |
+--------------------------+--------------------+------------------------------+
| Identifier               | 0x3                | SFP transceiver              |
| Ext. Identifier          | 0x4                | SFP/GBIC                     |
| Connector                | 0x7                | LC                           |
+--------------------------+--------------------+------------------------------+
| Transceiver Codes        | 0x010d001202000000 |                              |
|  - 10 GbE Compliance     | 0x00               |                              |
|  - SONET Compliance      | 0x0000             |                              |
|  - Ethernet Compliance   | 0x02               | 1000BASE-LX                  |
|  - Link Length           | 0x12               | Long distance (L)            |
|  - Transmitter Technology| 0x0012             | Longwave laser (LL)          |
|  - Transmission Media    | 0x0d               | Single Mode (SM)             |
|  - Channel speed         | 0x01               | 100 MBytes/Sec               |
+--------------------------+--------------------+------------------------------+
| Encoding                 | 0x01               | 8B10B                        |
| BR, Nominal              | 13                 | Gigabit                      |
|--------------------------+--------------------+------------------------------|
| Length(9um fiber) 1km    | 10                 | 10km                         |
| Length(9um fiber) 100m   | 100                | 10000m                       |
| Length(50um) 10m         | 55                 | 550m                         |
| Length(62.5um) 10m       | 55                 | 550m                         |
| Length(copper) 1m        | 0                  | 0m                           |
|--------------------------+--------------------+------------------------------|
| Vendor Name              | CIENA              |                              |
| Vendor OUI               | 0x000000           |                              |
| Vendor PN                | XCVR-010M31-03     |                              |
| Vendor Revision          | 10                 |                              |
| Vendor Serial Number     | B2R2011340         |                              |
| Vendor CLEI Code         | COUIA0SPAA         |                              |
| Ciena                    | XCVR-A00G85        |                              |
| Ciena Revision           | D                  |                              |
| Wavelength               | 0                  |                              |
|--------------------------+--------------------+------------------------------|
| Options                  | 0x1a               |                              |
|  - RATE_SELECT           | Bit 5              | No                           |
|  - TX_DISABLE            | Bit 4              | Yes                          |
|  - TX_FAULT              | Bit 3              | Yes                          |
|  - Loss of Signal Invert | Bit 2              | No                           |
|  - Loss of Signal        | Bit 1              | Yes                          |
|--------------------------+--------------------+------------------------------|
| BR, max                  | 0                  |                              |
| BR, min                  | 0                  |                              |
| Vendor Serial Number     | A9640060800547     |                              |
| Date (mm/dd/yy)          | 03/15/06           |                              |
|--------------------------+--------------------+------------------------------|
```

Diag Monitor Type	0x0	
- Legacy diagnostics	Bit 7	No
- Diagnostics monitoring	Bit 6	No
- Internally calibrated	Bit 5	No
- Externally calibrated	Bit 4	No
- Rx power measurement	Bit 3	OAM
Enhanced Options	0x0	
- Alarm/Warning Flags	Bit 7	No
- Soft TX_DISABLE	Bit 6	No
- Soft TX_FAULT	Bit 5	No
- Soft RX_LOS	Bit 4	No
- Soft RATE_SELECT	Bit 3	No
SFF-8472 Compliance	0x0	None

Procedure 5-16
Determining transceiver speed

When a transceiver is plugged in, the port speed is blank until a link is established, and then it is set to match the transceiver speed.

The Encoding “Value” column displays the actual value read from the optic, while the “Decoded String Equivalent” column indicates the supported port speed.

Step	Action

1 Display transceiver information:

```
port xcvr show
```

2 Display transceiver vendor data:

```
port xcvr show port <port> vendor
```

where

port <port> is the port.

vendor displays transceiver vendor data.

—end—

Example

This example shows sample output for the port show command.

```
> port show
```

```
+-------------------------------------------------------------------------------+
| Port Table      |           Operational Status             |  Admin Config    |
|-----------------+------------------------------------------+------------------|
| Port | Port     |      | Link State    |      |     |      | Auto |      |      | Auto |
| Name | Type     | Link |   Duration    | XCVR | STP | Mode | Neg  | Link | Mode | Neg  |
|------+----------+------+---------------+------+-----+------+------+------+------+------|
| 1    |10/100/G  | Up   |   0d 5h 4m27s |      | FWD | 100/FD | On | Ena  |1000/FD | On |
| 2    |10/100/G  | Up   |   0d 5h 4m27s |      | FWD | 100/FD | On | Ena  |1000/FD | On |
| 3    |10/100/G  |Down  |   0d 0h 0m 0s |      | Dis |      | On | Ena  |1000/FD | On |
| 4    |10/100/G  |Down  |   0d 0h 0m 0s |      | Dis |      | On | Ena  |1000/FD | On |
| 5    |10/100/G  |Down  |   0d 0h 0m 0s |      | Dis |      | On | Ena  |1000/FD | On |
| 6    |10/100/G  |Down  |   0d 0h 0m 0s |      | Dis |      | On | Ena  |1000/FD | On |
| 7    |10/100/G  |Down  |   0d 0h 0m 0s |      | Dis |      | On | Ena  |1000/FD | On |
| 8    |10/100/G  |Down  |   0d 0h 0m 0s |      | Dis |      | On | Ena  |1000/FD | On |
| 9    |10/100/G  |Down  |   0d 0h 0m 0s |      | Dis |      | On | Ena  |Auto/FD | On |
| 10   |10/100/G  |Down  |   0d 0h 0m 0s |      | Dis |      | On | Ena  |Auto/FD | On |
| 11   |10/100/G  |Down  |   0d 0h 0m 0s |      | Dis |      | On | Ena  |Auto/FD | On |
| 12   | Gig      | Up   |   4d 2h48m52s | Ena  | FWD | 1000/FD | On | Ena  |Auto/FD | On |
+------+----------+------+---------------+------+-----+------+------+------+------+------+
```

This example shows sample output for the port xcvr show port <PortNameList> vendor command.

```
> port xcvr show port 1 vendor

+------------------------ XCVR VENDOR DATA - Port 1 ----------------------------+
| Parameter                | Value              | Decoded String Equivalent     |
+--------------------------+--------------------+-------------------------------+
| Identifier               | 0x3                | SFP transciever               |
| Ext. Identifier          | 0x4                | SFP/GBIC                      |
| Connector                | 0x7                | LC                            |
+--------------------------+--------------------+-------------------------------+
| Transceiver Codes        | 0x0000000002000000 |                               |
|  - SONET Compliance      | 0x0000             |                               |
|  - Ethernet Compliance   | 0x02               | 1000BASE-LX                   |
|  - Link Length           | 0x00               | unknown                       |
|  - Transmitter Technology| 0x0000             | unknown                       |
|  - Transmission Media    | 0x00               | unknown                       |
|  - Channel speed         | 0x00               | unknown                       |
+--------------------------+--------------------+-------------------------------+
| Encoding                 | 0x01               | 8B10B                         |
| BR, Nominal              | 13                 | Gigabit                       |
|--------------------------+--------------------+-------------------------------|
...
+--------------------------+--------------------+-------------------------------+
```

Procedure 5-17
Tuning XFP transceivers

Tunable XFPs provide the ability to set the laser frequency in gigahertz (GHz), wavelength in nanometers with decimals, or set the channel. When you set the value for the frequency, wavelength, or channel, the other parameters are automatically populated.

If the range is set out of the supported range for the XFP, an error message is returned with the correct range. Attempting to set the tuning parameters for an SFP that does not support tunability generates an error.

Note: Tuning an XCVR causes a traffic outage lasting no more than a few seconds.

Step	Action

To tune XFP transceivers

1 View the current values for frequency, wavelength, or channel:

```
port xcvr show port <port> tunability
```

where

port <port>	is the port.
tunability	displays transceiver tunable data.

2 Set the frequency, wavelength, or channel:

```
port xcvr set port <port> {[frequency <NUMBER>] |
[wavelength <String>] | [channel <NUMBER>] [forward-
error-correction-mode <gfec | efec>]}
```

where

port <port>	is the port or ports to set.
frequency <NUMBER>	is the transceiver frequency in GHz.
wavelength <String>	is the transceiver wavelength in nanometers.
channel <NUMBER>	is the transceiver channel number.

To unset transceiver values

3 Unset transceiver values:

```
port xcvr unset port <port> {frequency} {wavelength}
{channel}
```

where

port <port>	is the port or ports to unset.
frequency <NUMBER>	is the transceiver frequency.
wavelength <String>	is the transceiver wavelength.
channel <NUMBER>	is the transceiver channel number.

—end—

Example

This example shows sample output for a port with a tunable XFP.

```
> port xcvr show port 2.1 tunability

+-------------------- XCVR Tunability - Port 2.1      -----------------+
| Field                |                      Value                     |
+-----------------------------------------------------------------------+
| Frequency Tunable    |                        Yes                     |
+----------------------+-----------------+-----------------+-----------+
|                      |          GHz    |            nm   |   ch#     |
|                      +-----------------+-----------------+-----------+
| Admin                |       191100    |       1568.8    |     1     |
| Oper Min             |       191100    |       1568.8    |     1     |
| Oper Max             |       196150    |       1528.4    |   102     |
| Oper Value           |       191100    |       1568.8    |     1     |
| Oper Error           |            0    |          0.0    |     0     |
+----------------------+-----------------+-----------------+-----------+
| Oper Grid Spacing    |                      50                       |
+----------------------+-----------------------------------------------+
```

This example shows sample output for a port without support for tunability.

```
> port xcvr show port 48 tunability

Not Supported for port 48
```

This example sets the frequency to 196150 on port 9.

```
> port xcvr set port 9 frequency 196150
```

Procedure 5-18
Tuning OTN FEC SFP+ transceivers

Tunable OTN FEC SFP+ transceivers (XCVR-TFEC01) have forward error correction (FEC) capability and configurable transceiver FEC mode (GFEC or EFEC) support.

Supported transceiver FEC modes are described in "Forward Error Correction mode configuration support" on page 5-17.

Step	Action

To tune OTN FEC SFP+ transceivers

1 Set the FEC mode:

```
port xcvr set port <port> [forward-error-correction-mode
<gfec | efec>]
```

where

port <port>	is the port or ports to set.
forward-error-correction-mode <gfec | efec>	is the forward error correction (FEC) mode of the transceiver.

To unset transceiver values

2 Unset transceiver values:

```
port xcvr unset port <port> [forward-error-correction-
mode]
```

where

port <port>	is the port or ports to unset.
forward-error-correction-mode	is the FEC mode of the transceiver.

—end—

Example

This example shows the configured FEC mode (GFEC/EFEC) by user:

```
3942*> port xcvr show port 21

+-------------------------- XCVR VENDOR DATA - Port 21   --------------------------+
| Parameter                | Value              | Decoded String Equivalent          |
+--------------------------+--------------------+------------------------------------+
| Identifier               | 0x3                | SFP transceiver                    |
| Ext. Identifier          | 0x4                | SFP/GBIC                           |
| Connector                | 0x7                | LC                                 |
+--------------------------+--------------------+------------------------------------+
| Transceiver Codes        | 0x0000000000000080 |                                    |
```

```
|  - 10 GbE Compliance       | 0x80               | 10GBASE-ER                        |
|  - SONET Compliance        | 0x0000             |                                   |
|  - Ethernet Compliance     | 0x00               | unknown                           |
|  - Link Length             | 0x00               | unknown                           |
|  - Transmitter Technology  | 0x0000             | unknown                           |
|  - Transmission Media      | 0x00               | unknown                           |
|  - Channel speed           | 0x00               | unknown                           |
+----------------------------+--------------------+-----------------------------------+
| Encoding                   | 0x06               | reserved                          |
| BR, Nominal                | 103                | Gigabit                           |
|----------------------------+--------------------+-----------------------------------|
| Length(9um fiber) 1km      | 80                 | 80km                              |
| Length(9um fiber) 100m     | 255                | 25500m                            |
| Length(50um) 10m           | 0                  | 0m                                |
| Length(62.5um) 10m         | 0                  | 0m                                |
| Length(copper) 1m          | 0                  | 0m                                |
|----------------------------+--------------------+-----------------------------------|
| Vendor Name                | CIENA-MEN          |                                   |
| Vendor OUI                 | 0x000000           |                                   |
| Vendor PN                  | XCVR-TFEC01        |                                   |
| Vendor Revision            | A                  |                                   |
| Vendor Serial Number       | 17230060           |                                   |
| Vendor CLEI Code           | CMUIASJGAA         |                                   |
| Ciena PN                   | XCVR-TFEC01        |                                   |
| Ciena Revision             | A                  |                                   |
| Wavelength                 | 1529.16            |                                   |
|----------------------------+--------------------+-----------------------------------|
| Options                    | 0x45a              |                                   |
|  - Tunable                 | Bit 6              | Yes                               |
|  - RATE_SELECT             | Bit 5              | No                                |
|  - TX_DISABLE              | Bit 4              | Yes                               |
|  - TX_FAULT                | Bit 3              | Yes                               |
|  - Loss of Signal Invert   | Bit 2              | No                                |
|  - Loss of Signal          | Bit 1              | Yes                               |
|----------------------------+--------------------+-----------------------------------|
| BR, max                    | 0                  |                                   |
| BR, min                    | 0                  |                                   |
| Date (mm/dd/yy)            | 07/27/17           |                                   |
|----------------------------+--------------------+-----------------------------------|
| Diag Monitor Type          | 0x68               |                                   |
|  - Legacy diagnostics      | Bit 7              | No                                |
|  - Diagnostics monitoring  | Bit 6              | Yes                               |
|  - Internally calibrated   | Bit 5              | Yes                               |
|  - Externally calibrated   | Bit 4              | No                                |
|  - Rx power measurement    | Bit 3              | Avg                               |
|----------------------------+--------------------+-----------------------------------|
| Enhanced Options           | 0xf0               |                                   |
|  - Alarm/Warning Flags     | Bit 7              | Yes                               |
|  - Soft TX_DISABLE         | Bit 6              | Yes                               |
|  - Soft TX_FAULT           | Bit 5              | Yes                               |
|  - Soft RX_LOS             | Bit 4              | Yes                               |
|  - Soft RATE_SELECT        | Bit 3              | No                                |
|----------------------------+--------------------+-----------------------------------|
| SFF-8472 Compliance        | 0x5                | Other                             |
+----------------------------+--------------------+-----------------------------------+

+---------------------- XCVR DIAGNOSTICS - Port 21       -----------------+
|                  |           |           Alarm          |        Warning        |
| Output           | Value     | Threshold      | Flag    | Threshold      | Flag |
+------------------+-----------+----------------+---------+----------------+------+
| Temp    (degC)   |   59.783  | HIGH  70.000   | 0       | HIGH  65.000   | 0    |
|                  |           | LOW    0.000   | 0       | LOW    5.000   | 0    |
+------------------+-----------+----------------+---------+----------------+------+
| Vcc     (volts)  |    3.220  | HIGH   3.560   | 0       | HIGH   3.460   | 0    |
|                  |           | LOW    3.040   | 0       | LOW    3.140   | 0    |
+------------------+-----------+----------------+---------+----------------+------+
| Bias      (mA)   |   50.912  | HIGH 130.000   | 0       | HIGH 120.000   | 0    |
|                  |           | LOW    5.000   | 0       | LOW   20.000   | 0    |
+------------------+-----------+----------------+---------+----------------+------+
| Tx Power  (mW)   |    1.084  | HIGH   2.511   | 0       | HIGH   1.990   | 0    |
```

```
|                |          | LOW    0.560   | 0    | LOW    0.707   | 0    |
+----------------+----------+----------------+------+----------------+------+
| Tx Power (dBm) |  +0.3531 | HIGH  +3.9985  | 0    | HIGH  +2.9885  | 0    |
|                |          | LOW   -2.5181  | 0    | LOW   -1.5058  | 0    |
+----------------+----------+----------------+------+----------------+------+
| Rx Power  (mW) |   0.0858 | HIGH   0.3988  | 0    | HIGH   0.3160  | 0    |
|                |          | LOW    0.0006  | 0    | LOW    0.0010  | 0    |
+----------------+----------+----------------+------+----------------+------+
| Rx Power (dBm) | -10.6651 | HIGH  -4.0012  | 0    | HIGH  -5.0031  | 0    |
|                |          | LOW  -32.2185  | 0    | LOW  -30.0000  | 0    |
+----------------+----------+----------------+------+----------------+------+
|                          Module Capabilites                              |
+------------------------+------------------------+-----------------------+
| Laser First Freq (GHz) | Laser Last Freq (GHz)  | Grid Spacing (GHz)    |
+------------------------+------------------------+-----------------------+
|     191500             |          196050        |            50         |
+------------------------+------------------------+-----------------------+
|                    Frequency and Wavelength Control                      |
+---------------------------+-----------------------------+---------------+
| Channel Number            | Wavelength (nm)             | TX Dither     |
+---------------------------+-----------------------------+---------------+
|         42                |  1548.9                     | 0             |
+---------------------------+-----------------------------+---------------+
|                    Frequency and Wavelength Errors                       |
+-------------------------------+-------------------------------------------+
| Frequency (GHz)               | Wavelength (nm)                           |
+-------------------------------+-------------------------------------------+
|            0                  |       0.00                                |
+-------------------------------+-------------------------------------------+
|                            Current Status                                 |
+--------+-------+--------+--------+--------+--------+--------+-----------+
|        | TEC   | WaveL  | Tx     |        |        |        |           |
|        | Fault | Unlock | Tune   |        |        |        |           |
+--------+-------+--------+--------+--------+--------+--------+-----------+
|   0    |   0   |    0   |   0    |   0    |   0    |   0    |   0       |
+--------+-------+--------+--------+--------+--------+--------+-----------+
|                            Latched Status                                 |
+--------+-------+--------+--------+--------+--------+--------+-----------+
|        | TEC   | WaveL  | Bad    | New    |Unsup Tx|        |           |
|        | Fault | Unlock |Channel |Chnnel  | Dither |        |           |
+--------+-------+--------+--------+--------+--------+--------+-----------+
|   0    |   0   |    0   |   0    |   0    |   0    |   0    |   0       |
+--------+-------+--------+--------+--------+--------+--------+-----------+
|             OTN Forward Error Correction(FEC) Capabilities                |
+-----------------------------------+-----------------------------------+
| Status                            | Enabled                           |
+-----------------------------------+-----------------------------------+
| Mode                              | Enhanced(EFEC)                    |
+-----------------------------------+-----------------------------------+
```

Procedure 5-19
Setting the port connector mode

Some ports (called dual mode or combination ports) support both RJ45 and SFP (including smaller size SFP+) connectors. Only one of these connectors can be active at a given time. Some ports support a default connector mode, where the mode operates as SFP if a transceiver is installed, or an RJ45 if not. You can set the connector mode manually for the port. This table shows the dual mode ports and default mode for each platform that supports them.

Table 5-5
Factory default general port settings by platform

Platform	Ports	Default Connector Mode
3903	1	Default
3903x	1	Default
3904	1-2	Default
3905	1-2	Default
3906mvi	1-4	Default
3916	1-2	Default
3930	1-4	Default
3932	1-4	Default

Note: The port connector mode is applicable only to combination ports that support RJ45 or SFP connectors. It does not apply to XFP ports.

In addition, speed is set to “Auto” with auto-negotiation enabled. So, you can install 1G or 100M transceivers, and the system automatically set the speed accordingly. If these settings or other port attributes are set explicitly and do not match the capabilities of the active connector, a mismatch warning is generated. The warning is cleared when the attributes match the capabilities of the active connector.

Note: If you attempt to set the mode to a connector that is not supported for the specified port, the system generates a “Capability not supported” error message.

Step	Action

1 Set the mode for a specific port:

```
port set port <port> mode <default|rj45|sfp>
```

where

port <port>	is the port.
mode <default|rj45|sfp>	is the physical interface connector mode.

—end—

Example

This example sets the mode for port 9 to RJ45.

```
port set port 9 mode rj45
```

Link aggregation

Link aggregation is defined in IEEE 802.1AX. This standard defines how two or more full-duplex Ethernet ports of the same speed can be combined into a single logical port to carry traffic between two devices connected in parallel. This logical grouping of ports enables load sharing and load balancing among these ports and thus an aggregation of bandwidth as well. Traffic destined to egress on an aggregated port is distributed among all the links in the group.

Link aggregation also provides inherent, automatic redundancy for high-traffic network connections. This is achieved by dynamically redirecting traffic from a failed port to the remaining good ports in the aggregation group. Link aggregation can be used to expand bandwidth and add link redundancy.

> ***Note 1:*** The IEEE 802.1AX standard does not support aggregation of bandwidth between links taking different paths and connecting to different upstream units. Only ports connected in parallel between two units can be aggregated.
>
> ***Note 2:*** You cannot create an aggregation that uses one non-enhanced and one enhanced port on the 5150.

For the related procedures, see “Configuring LACP between two devices” on page 6-26 and “Configuring LACP protection” on page 6-29.

See *39xx/51xx Service Delivery, Aggregation and Virtualization Command Reference* for additional aggregation commands.

Link aggregation groups

Link aggregation is configured using the concept of physical ports vs. logical ports. Physical ports are actual ports on the switch, while a logical port or Link Aggregation Group (LAG) is the collective group the physical ports belong to. Grouped physical ports appear to the system as one logical port. For example, with spanning tree, a physical port assigned to a LAG appears to be operationally down. This excludes the physical port from being considered in the spanning tree topology.

In the LAG group there is always a ‘lead port’, which is responsible for sending control frames for protocols such as RSTP. The lowest port number in the LAG is always elected as the lead port. If the lead port goes down, the port with the next lowest port number assumes the role of lead port. If a link failure occurs, it is almost transparent. The only traffic losses are the frames either in the egress queue or those already on the link at the time of failure.

Link aggregation supports these hashing mechanisms:

- Known Unicast: layer 2, layer 3 and enhanced (3903, 3903x, 3904, 3905, 3906mvi, 3916, 3926m, 3928, 3930, 3931, 3932, 3942, 5142, 5150, and 5160 platforms)
- Unicast, multicast, broadcast (BUM): simplified and enhanced (3903, 3903x, 3904, 3905, 3906mvi, 3916, 3926m, 3928, 3930, 3931, 3932, 3942, 5142, 5150, and 5160 platforms) flood hash modes

This table shows the number of LAGs supported for each platform, and whether enhanced hashing is supported.

Table 6-1
Supported LAGs per platform

Platform	Number of LAGs	Number of physical ports per LAG	Enhanced Hashing Support
3903	3	3	Yes
3903x	3	3	Yes
3904	4	4	Yes
3905	4	4	Yes
3906mvi	6	6	Yes
3916	6	6	Yes
3926m	8	8 (4 primary + 4 protection) (see Note on page 6-3)	Yes
3928	12	8 (4 primary + 4 protection) (see Note on page 6-3)	Yes
3930	10	8	Yes
3931	10	8	Yes
3932	10	8	Yes
3942	24	8	Yes

Table 6-1
Supported LAGs per platform

Platform	Number of LAGs	Number of physical ports per LAG	Enhanced Hashing Support
5142	24	8	Yes
5150	52	8	Yes
5160	24	8	Yes
Note: This device supports a maximum of four member ports in LAG distribution. This implies that up to four member ports can be added in Active:Active LAG. In the case of protection LAG, up to four primary and four protection ports can be added.			

In addition aggregated ports operate as if they were a single entity with respect to L2 address learning, regardless of the port on which the LAG traffic ingresses.

Two types of Link Aggregation are supported, manual Link Aggregation and Link Aggregation Control Protocol (LACP). Regardless of the type of Link Aggregation, the user is only required to create a LAG, add ports to it, and then specify the mode of aggregation (manual or LACP). The device takes care of the rest of the configuration itself.

Devices support the creation of as many aggregation groups as number of physical ports on the device. This lets the user to always provision services over aggregation ports so that services can be easily moved to a different port.

The hashing mode for known unicast frames is configurable per aggregation, and can be changed any time during the life of the aggregation. Note that the hashing algorithms are probabilistic and do not in any way guarantee that traffic is distributed equally among all the distributing ports of the aggregation. The hashing mode for unknown unicast, multicast and broadcast traffic is configurable on a per-device-basis.

Hashing mechanisms

There are several configuration options available to specify the type of hashing algorithm to use based on the frame type. There are two options:

- Known unicast (MAC address lookup) supporting layer 2, layer 3 and enhanced hashing
- Unicast, multicast, broadcast (BUM) supporting simplified and enhanced flood hashing

Link aggregation supports these hashing mechanisms:

- Known unicast: layer 2, layer 3 and enhanced (3903, 3903x, 3904, 3905, 3906mvi, 3916, 3926m, 3928, 3930, 3931, 3932, 3942, 5142, 5150, and 5160 platforms)
- Unicast, multicast, broadcast (BUM): simplified and enhanced (3903, 3903x, 3904, 3905, 3906mvi, 3916, 3926m, 3928, 3930, 3931, 3932, 3942, 5142, 5150, and 5160 platforms)

The hashing mechanism is common across link aggregation groups.

Hashing keys depend on the type of traffic and the selected hash mode, as shown in this table.

Table 6-2
Hashing keys

Traffic type	Hash mode	Flood Hash	Hashing keys
Layer 2 Known Unicast	MAC based address	Any	• Source port • Destination MAC • Source MAC • VID • Ethertype
Layer 2 Unknown Unicast, Multicast and Broadcast	Any	Simplified	• Source port • Destination MAC • Source MAC
	Any	Enhanced	• Source port • Destination MAC • Source MAC • VID [1] • EtherType

Table 6-2
Hashing keys (continued)

Traffic type	Hash mode	Flood Hash	Hashing keys
IPv4 Known Unicast	MAC address based	Any	• Source port • Destination MAC • Source MAC • VID • EtherType
	IP address based	Any	• Source IP address • Destination IP address • TCP/UDP source port • TCP/UDP destination port
	Enhanced	Any	• Source IP address • Destination IP address • IP protocol • TCP/UDP source port • TCP/UDP destination port
IPv4 Unknown Unicast, Multicast and Broadcast	Any	Simplified	• Source port • Destination MAC • Source MAC
	Any	Enhanced	• Source IP address • Destination IP address • IP protocol • TCP/UDP source port • TCP/UDP destination port

Table 6-2
Hashing keys (continued)

Traffic type	Hash mode	Flood Hash	Hashing keys
IPv6 Known Unicast	MAC address based	Any	• Source port • Destination MAC • Source MAC • VID • EtherType
	IP address based	Any	• Source IPv6 address • Destination IPv6 address • TCP/UDP source port • TCP/UDP destination port
	Enhanced	Any	• Source IPv6 address • Destination IPv6 address • Flow Label [2] • Next Header • TCP/UDP source port • TCP/UDP destination port
IPv6 Unknown Unicast, Multicast and Broadcast	Any	Simplified	• Source port • Destination MAC • Source MAC
	Any	Enhanced	• Source IPv6 address • Destination IPv6 address • Flow Label • Next Header • TCP/UDP source port • TCP/UDP destination port

Table 6-2
Hashing keys (continued)

Traffic type	Hash mode	Flood Hash	Hashing keys
Terminating MPLS Known Unicast: 3903, 3903x, 3904, 3905, 3906mvi, 3916, 3930, 3931, 3932, 5150 — Any L2 payload 3942, 5142 and 5160 — Non-IP payload	MAC addressed based	Any	• Source port • Tunnel destination MAC • Tunnel source MAC • Tunnel VID • EtherType
	IP address based	Any	• Source port • Tunnel destination MAC • Tunnel source MAC • Tunnel VID • EtherType
	Enhanced	Any	• Source port • Customer source MAC address • Customer destination MAC address • Customer Ethertype
Terminating MPLS Unknown Unicast, Multicast and Broadcast: 3903, 3903x, 3904, 3905, 3906mvi, 3916, 3930, 3931, 3932, 5150 — Any L2 payload 3942, 5142 and 5160 — Non-IP payload	Any	Simplified	• Source port • Tunnel destination MAC • Tunnel source MAC
	Any	Enhanced	• Source port • Customer source MAC address • Customer destination MAC address • Customer EtherType

Table 6-2
Hashing keys (continued)

Traffic type	Hash mode	Flood Hash	Hashing keys
Terminating MPLS Layer 2 Known Unicast with IPv4 payload (3942, 5142 and 5160 devices)	MAC address based	Any	• Source port • Tunnel destination MAC address • Tunnel source MAC address • Tunnel VID • EtherType
	IP addressed based	Any	• Source port • Tunnel destination MAC address • Tunnel source MAC address • Tunnel VID • EtherType
	Enhanced	Any	• Source port • Customer source IP address • Customer destination IP address • Customer IP protocol • Customer L4 source port • Customer L4 destination port
Terminating MPLS Layer 2 Unknown Unicast, Multicast and Broadcast with IPv payload (3942, 5142 and 5160 devices)	Any	Simplified	• Source port • Tunnel destination MAC address • Tunnel source MAC address
	Any	Enhanced	• Source port • Customer source IP address • Customer destination IP address • Customer IP protocol • Customer L4 source port • Customer L4 destination port

Table 6-2
Hashing keys (continued)

Traffic type	Hash mode	Flood Hash	Hashing keys
Terminating MPLS Layer 2 Known Unicast with IPv6 payload (3942, 5142 and 5160 devices)	MAC address based	Any	• Source port • Tunnel destination MAC address • Tunnel source MAC address • Tunnel VID • EtherType
	IP addressed based	Any	• Source port • Tunnel destination MAC address • Tunnel source MAC address • Tunnel VID • EtherType
	Enhanced	Any	• Source port • Customer destination IPv6 • Customer source IPv6 • Customer L4 source port • Customer L4 destination port • Customer flow label • Customer next header
Terminating MPLS Layer 2 Unknown Unicast, Multicast and Broadcast with IPv6 payload (3942, 5142 and 5160 devices)	Any	Simplified	• Source port • Tunnel destination MAC address • Tunnel source MAC address
	Any	Enhanced	• Source port • Customer destination IPv6 • Source IPv6 • Customer L4 source port • Customer L4 destination port • Customer flow label • Customer next header

Table 6-2
Hashing keys (continued)

Traffic type	Hash mode	Flood Hash	Hashing keys
MPLS Non-terminating Known Unicast	MAC based address	Any	• Source port • Destination MAC address • Source MAC address • VID • EtherType
	IP based address	Any	• Source port • Destination MAC address • Source MAC address • VID • EtherType
	Enhanced	Any	• Source port • Top MPLS label • Second MPLS label • Third MPLS label (3942, 5142 and 5160 devices only)
MPLS Non-terminating Unknown Unicast, Multicast and Broadcast	Any	Simplified	• Source port • Destination MAC address • Source MAC address
	Any	Enhanced	• Source port • Top MPLS label • Second MPLS label • Third MPLS label (3942, 5142 and 5160 devices only)

1.The VID field in either enhanced mode is the frame's outer VID as it was received on the ingress port.

2.IPv6 Flow Label is used as a hash input only on the 3942, 5142 and 5160 platforms.

3.L4 source port and L4 destination port refers to port numbers in TCP and UDP headers only.

Note: EOAM loopback set on the LEAD port of an aggregation causes the RECEIVE-side aggregation (if it is the lead port in loopback) to NOT forward traffic. EOAM loopback done on ANY other physical port in the aggregation must pass traffic correctly.

Known unicast traffic hashing

There are three LAG hash modes available for distributing known unicast frames:

- Layer 2 MAC address-based hashing
- Layer 3 IP address-based hashing
- Enhanced hashing, available on the 3903, 3903x, 3904, 3905, 3906mvi, 3916, 3926m, 3928, 3930, 3931, 3932, 3942, 5142, 5150 and 5160 platforms

The hash mode used for unicast frames is chosen through configuration on a per-aggregation port basis. The 39XX/51XX devices use a Layer 2 address-based hashing mechanism by default. Table 6-2 shows the fields used to compute the hash for known unicast frames.

Layer 2 MAC address based hashing uses these fields of an Ethernet frame to compute hash values for known unicast frames:

- Source port
- Destination MAC address
- Source MAC address
- VLAN (if present)
- EtherType

Layer 3 hashing uses these fields in a packet to compute a hash value for known unicast frames that contain either an IPv4 or IPv6 EtherType value:

- Source IP address
- Source TCP/UDP port
- Destination IP address
- TCP/UDP destination port

The aggregate is configured with an IP or L3 address-based hash mode, but frames that do not contain an IPv4 or IPv6 EtherType value are processed using the applicable layer 2 hash algorithm.

Enhanced hashing provides greater flexibility on the hash algorithm used based on the known unicast frame type. The difference between the Enhanced hashing mechanism and the Layer 2 and Layer 3 hashing mechanisms for known unicast frames include:

- IPv4 hash algorithm includes the IP Protocol field.

- IPv6 hash algorithm includes the Next Header field and the Flow Label (3942, 5142 and 5160 platforms only) field.
- Most of the bits of all hash input fields are used in the enhanced hash calculation, whereas the other hash modes on the 3903, 3903x, 3904, 3905, 3906mvi, 3916, 3926m, 3928, 3930, 3931, 3932, 5150 and 5160 platforms only consider the three least significant bits of each byte in the input fields.
- All hash algorithms result in an eight-bit hash value versus a three-bit hash value on the 3903, 3903x, 3904, 3905, 3906mvi, 3916, 3926m, 3928, 3930, 3931, 3932, 3942, 5150 and 5160 platforms.
- LAG member selection is performed using a modulo calculation on the eight-bit hash result. This distributes frames on 256-hash bins providing better traffic distribution across LAG members.
- Supports specific hash algorithms for MPLS terminating and non-terminating frames.

Unicast, multicast and broadcast traffic hashing

Simplified hash mode and enhanced hash mode are used to distribute unknown unicast, multicast and broadcast (BUM) frames.

The simplified hash mode is supported on all devices, but the configurable attribute is only supported on the devices that support the enhanced hash mode. Simplified hash mode and enhanced has mode are supported on the 3903, 3903x, 3904, 3905, 3906mvi, 3916, 3926m, 3928, 3930, 3931, 3932, 3942, 5142, 5150 and 5160 platforms. The mode is chosen through the configuration on a per-device basis, with simplified mode as the system default.

The simplified default option uses the same algorithm used for multicast and broadcast frames and both L2 and L3 unknown unicast frames. The algorithm used is device dependent.

Enhanced hashing provides greater flexibility on the hash algorithm used for unknown unicast, multicast and broadcast frames. The differences between the enhanced hashing mechanism and the simplified hashing mechanism for BUM frames include:

- Supports specific hash algorithms for more specific frame types, including IPv4 frames and IPv6 frames
- Supports MPLS non-terminating and terminating frames.
- All the bits of all hash input fields are used in the enhanced hash calculation. In the simplified mode, the four least significant bits of input fields on 3903, 3903x, 3904, 3905, 3906mvi, 3926m, 3928, 3916, 3930, 3931, 3932, and 5150 platforms.

Software downgrade

For a software downgrade to a release that does not support the enhanced hash mode, aggregations that are assigned the enhanced hash mode are replaced by the IP address based hash mode. Although the IP address based hash mode is not the default hash mode, it has the hashing capabilities that are closest to enhanced, particularly for known unicast IP frames.

Manual link aggregation

LACP allows the device to negotiate link aggregation by sending LACP PDUs to its peer. Manually configured Link Aggregation ports are considered “blind” as they are set up completely by the administrator and never send LACP PDUs. A manual Link Aggregation forms only when a LAG has been specifically configured on two devices following which, their ports are physically connected.

Although Link Aggregation has been a standard since 1998 every vendor has a slightly different implementation (e.g., Cisco Ether Channel, Sun Trunking, etc.). Manual Link Aggregation is most useful when connecting a Ciena device to another vendor’s device that supports manual Link Aggregation.

LACP

IEEE 802.1AX LACP allows for standards-based link aggregation between devices over full duplex, same-speed links. LACP utilizes control packets to establish aggregation of links, maintain that aggregation, and distribute or redistribute bandwidth for load balancing and link failure. The LACP protocol is very intricate and prone to configuration errors. Ciena has simplified the configuration of Link Aggregation.

Traditionally, network administrators had to configure parameters such as Actor Admin Keys, Operational keys, and collector max delay. In the Ciena implementation, all of these parameters are handled behind the scenes. After a LAG has been created on one device and ports are added, it looks for eligible ports on its peer to aggregate with. It doesn’t matter if the two LAGs have the same configurations as their peer, as long as the groups are valid they find each other and negotiate an aggregation, as shown in this figure.

Figure 6-1
Link Aggregation

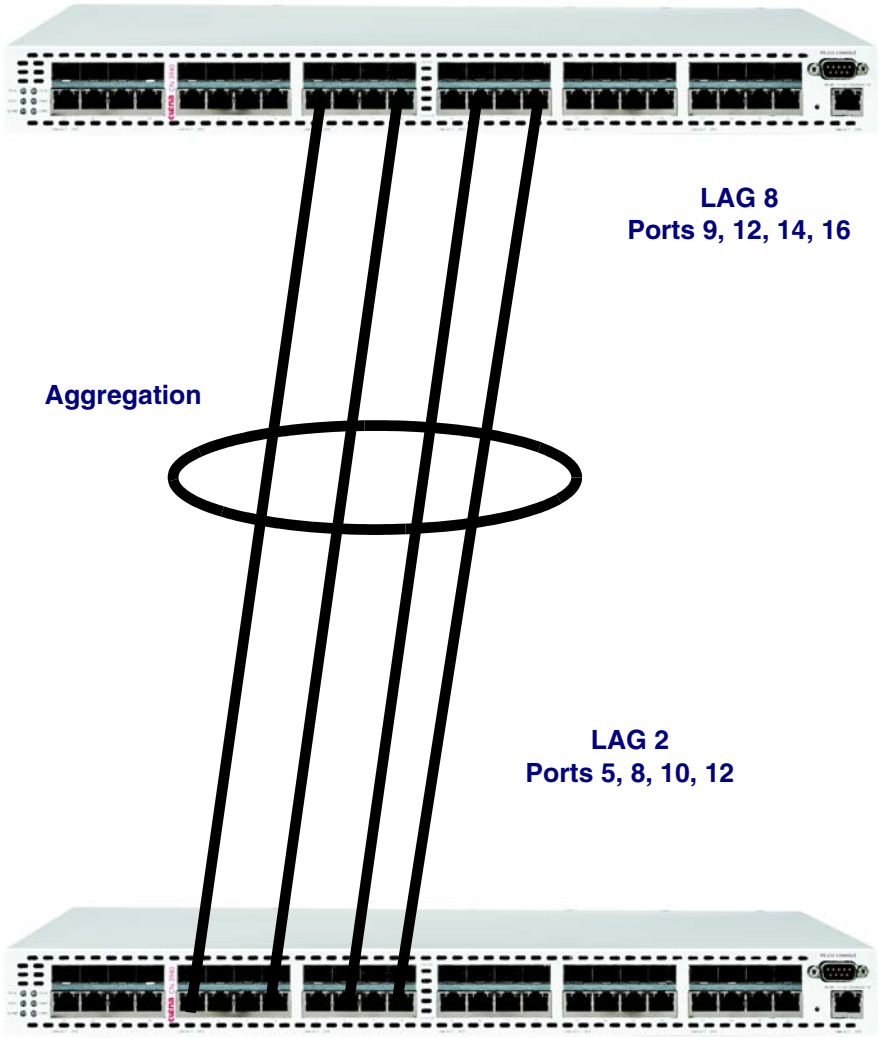

Note: Advanced users can still directly configure parameters and keys for both LAGs and physical ports.

LACP uses two types of control frames; LACP PDUs and Marker Frames. Both messages use a globally unique destination address of 01-80-C2-00-00-02 and a corresponding Ethernet type. LACP PDUs are used for negotiation between links. Marker frames are used to maintain frame sequence order.

As bandwidth is distributed across aggregated ports and a port is made ready to forward frames on the LAG, Marker Frames and Marker Response Frames are exchanged. When an event that requires flow redistribution occurs, the

flows that need to be redistributed are first blocked. Once blocked, a Marker Frame that has a response timer tied to it, is sent out each port in the aggregation group to the link partner for each flow. When the Marker Response Frames are received or the response timers expire, the flow is moved in the forwarding database of the newly selected port and forwarding is once again enabled.

Protection link aggregation

The Protection LAG feature allows the total number of member ports in a LAG to exceed the number of ports allowed in its distribution. Ports that are not actively distributing traffic are held in a standby state. When a distributing port becomes operationally disabled or deselected from the LAG, one of the standby ports replaces it in the distribution.

Each port is administratively designated as a Distribution or Protection port when it is added as a LAG member. Although this designation affects port precedence when active links are chosen, both Distribution and Protection ports can operate in either an active or a standby role.

For example, suppose that 3 ports are designated as Distribution ports, and 2 ports are designated as Protection ports. All 5 ports are grouped together in the same LAG. The Distribution ports handle all the traffic, and the Protection ports are essentially on standby. If a Distribution ports fails, then the Protection port with the higher priority takes its place.

By default, when a failed Distribution port returns to operational status, it is held in standby mode. This behavior can be overridden by enabling the revert-protection parameter. When the revert option is configured, along with the revert-delay timer, the original Distribution port becomes active again and the Protection port then reverts back to being a standby port. The timer delay protects the LAG from a port that may be rapidly flapping between enabled and disabled.

This table lists the roles that LAG ports can serve.

Table 6-3
LAG port roles

LAG port role	Description
Distribution Port	A port associated with a link that is intended to carry traffic. It has precedence over Protection ports when the LAG distribution (set of active links) is initially chosen. The number of member Distribution ports in a LAG determines the maximum number of ports allowed to distribute at one time.
Protection Port	A port associated with a link that is intended to act as a standby. It does not carry traffic unless a Distribution port is removed from the LAG distribution, in which case it transitions to an active role.
Active Port	A port, either Distribution or Protection, which is currently in the LAG distribution and is actively carrying traffic. For example, a Protection port becomes active when it replaces a Distribution port.
Standby Port	A port, either Distribution or Protection, that is a member of a LAG but is not currently distributing traffic. For example, a Distribution port can become a Standby port when it becomes operational after a link failure, and reversion is disabled.

Note: For optimal performance, LAG member ports must be connected in the same order at both LAG partner nodes (that is, LAG member ports must not be connected in crisscross order). Details are provided in this example:

Example

The following is correct:

```
DUT-1_P1 <====> DUT-2_P5
DUT-1_P2 <====> DUT-2_P11
DUT-1_P10<====> DUT-2_P12
```

The following is incorrect:

```
DUT-1_P6 <====> DUT-2_P7 /** here ports 6 and 7 are connected
DUT-1_P7 <====> DUT-2_P6     in crisscross order **/
DUT-1_P8 <====> DUT-2_P8
```

Proprietary VS standard protection mode

An LACP LAG can operate in one of two distinct operating modes: Proprietary or Standard. The protection mode determines the LACP state of Standby ports and whether or not synchronizing the active link selection between the Actor and Partner systems is driven by LACP or requires consistent configuration between the two systems.

Proprietary Protection mode is a Ciena LACP mode that relies on consistent configuration to synchronize the active link selection. In Proprietary protection mode, a Standby port remains attached to the LAG, but the Distributing bit for the Actor State information is set to False. Since a Standby port advertises that it is IN_SYNC and Collecting, the partner system is technically allowed to distribute traffic over the link with the expectation that the traffic is forwarded normally.

Standard Protection mode strictly adheres to the Standby link operation as defined in the LACP standard. The set of active links is fully communicated through LACP state information and the system with the highest configured priority in the LAG has control over the active port selection. In Standard protection mode, a Standby port does not attach to the LAG until it becomes Active. As a result, its Actor State indicates that it is OUT_OF_SYNC, and neither Collecting nor Distributing. A link must be synchronized on both systems before either system can Collect or Distribute over it.

Proprietary protection mode

When configuring protection link aggregation, the configuration must match at both ends of the LAG (for example, protection port or distribution port, priority, and revert parameters). The partner system must also support this mode to ensure that both systems always distribute over the same set of links. Active port selection is more directly influenced by system configuration. This mode is not supported on non-Ciena devices.

Both the actor and partner systems must be running in Proprietary protection mode with a consistent configuration, so that both systems always choose the same set of active links. LACP is unable to ensure that this is the case, since the decision to transition from the Collecting to Distributing states is local once the link partner indicates that it is also Collecting.

Standby port precedence is determined by Port Aggregation Priority, which is a combination of configured port priority and port number. The ports at both ends of a link must be configured with the same priority relative to the other member ports, or each system may activate a different Standby port after a link failure.

The Example of initial LAG state, proprietary mode table depicts the local LAG configuration and operational states of a LAG with two Distribution and two Protection ports operating in Proprietary protection mode. The Distribution ports are the Active ports, while the Protection ports are the Standby ports. Typically, the LAG initially comes up in this state. Standby ports are Attached, IN_SYNC and Collecting, but not Distributing.

Table 6-4
Example of initial LAG state, proprietary mode

LAG port configuration			Port operational state			Actor state			Partner state		
Port	**Type**	**Priority**	**Oper**	**Select**	**Attach**	**Synch**	**Coll**	**Dist**	**Sync**	**Coll**	**Dist**
1	Dist	0x8000	Enabled	Active	Yes	X	X	X	X	X	X
2	Dist	0x8000	Enabled	Active	Yes	X	X	X	X	X	X
3	Prot	0x8000	Enabled	Standby	Yes	X	X		X	X	
4	Prot	0x4000	Enabled	Standby	Yes	X	X		X	X	

The LAG state after distribution link failure, proprietary mode table depicts the LAG after a link failure causes port 1 to become operationally disabled. Since port 4 is configured with a higher priority (numerically lower priority than port 3), it becomes Active and replaces port 1 in the LAG distribution. To reflect this, the Distributing bit is set to True in its Actor State.

Port 1 remains attached to the LAG and its Actor State is still IN_SYNC, although its Partner State has changed to OUT_OF_SYNC. This state is in accordance with the LACP standard, and it allows existing Partner State information to be reserved for quick reattachment when the link is restored. The link is precluded from becoming active while the port is not operational because it is not fully synchronized.

Table 6-5
LAG state after distribution link failure, proprietary mode

LAG port configuration			Port operational state			Actor state			Partner state		
Port	**Type**	**Priority**	**Oper**	**Select**	**Attach**	**Synch**	**Coll**	**Dist**	**Sync**	**Coll**	**Dist**
1	Dist	0x8000	Disabled	Standby	X	X				X	X
2	Dist	0x8000	Enabled	Active	X	X	X	X	X	X	X
3	Prot	0x8000	Enabled	Standby	X	X	X		X	X	
4	Prot	0x4000	Enabled	Active	X	X	X	X	X	X	X

Remember that the behavior for reversion depends on how it is configured. For example, after port 1 returns to operationally enabled, since Distribution port 2 and Protection port 4 are currently the Active ports, Distribution port 1 must assume a Standby role. As a result, it is attached, IN_SYNC and Collecting, but not Distributing.

If reversion is disabled on the LAG, then the ports stay in this state until either port 2 or 4 becomes operationally disabled or deselected, allowing port 1 to resume its Active role. If reversion is enabled, then a reversion timer is started for each Distribution port when it returns to normal. When the reversion timer expires, the Distribution port is allowed to become Active, and forces the Protection port back into a Standby role. Reversion allows the LAG to eventually return to its initial configured state as depicted in the Example of initial LAG state, proprietary mode table.

Standard protection mode

In Standard protection mode, a Standby port does not attach to the LAG until it becomes Active. As a result, its Actor State indicates that it is OUT_OF_SYNC, and is not Collecting or Distributing. Since a link must be synchronized on both partner systems before either can collect or distribute over it, LACP ensures that the partner does not include a Standby port in its distribution.

For this protection scheme to work, one system must have control over the active port selection, while the other system must respond by distributing over the same set of ports. The LAG never stabilizes if the two systems attempt to place the same port in a different state. To avoid this conflict, the higher priority system always controls which ports are Active. The higher priority system is determined by which system has the numerically lower configured LAG System Priority.

When the actor system has a lower priority, all ports selected to the LAG are given an Active port role, allowing them to be attached to the LAG, and to be IN_SYNC. Since the partner system is controlling the LAG, only the ports that the partner selects to be active have a Partner State of IN_SYNC, while the ports that the partner selects to be Standby have a Partner State of OUT_OF_SYNC. Only ports that are IN_SYNC on both systems are able to join the LAG distribution on either system, thus the LACP state ensures that the systems are always using the same set of links.

In contrast, when the actor system has a higher priority, Standby ports remain in a waiting state, where they are selected but not yet attached to the LAG. Transitioning a Standby port into an active port is accomplished by completing the attachment, which in turn sets the port's Actor State to IN_SYNC. The port is already in an Active state on the lower priority partner system, so its Partner State is already IN_SYNC and the actor can immediately move the port to

Collecting and ultimately Distributing. The partner must receive an updated LACPDU indicating that the port is fully synchronized before it can initiate this process.

An additional consequence of Standard protection mode is that the actor port priority only has local significance. While actor port priority is still the most significant component of the Port Aggregation Priority value used to determine the precedence among Standby ports when the actor system has a higher priority, it has no influence when the Partner system is in control.

The Example of Initial LAG state, standard mode with actor priority table depicts the local LAG configuration and operational states of a LAG with two Distribution and two Protection ports operating in Standard protection mode, when the actor system has higher priority than the partner system. The Distribution ports are the Active ports, while the Protection ports are the Standby ports. Typically, the LAG initially comes up in this state. Standby ports are not Attached to the LAG, NOT_IN_SYNC, and neither Collecting, nor Distributing. Although a Standby port in Standard protection mode is not actually attached to the LAG, it is still selected to it, and does not function as an individual physical port while in this state.

Table 6-6
Example of Initial LAG state, standard mode with actor priority

LAG port configuration			Port operational state			Actor state			Partner state		
Port	**Type**	**Priority**	**Oper**	**Select**	**Attach**	**Synch**	**Coll**	**Dist**	**Sync**	**Coll**	**Dist**
1	Dist	0x8000	Enabled	Active	X	X	X	X	X	X	X
2	Dist	0x8000	Enabled	Active	X	X	X	X	X	X	X
3	Prot	0x8000	Enabled	Standby					X		
4	Prot	0x4000	Enabled	Standby					X		

The LAG state after distribution link failure, standard mode with actor priority table depicts the LAG after a link failure causes port 1 to become operationally disabled. The actor system is higher priority, allowing it to choose the replacement for port 1. Since port 4 is configured with a higher priority (numerically lower priority) than port 3 on the actor system, it becomes Active and replaces port 1 in the distribution. Transitioning port 4 from a Standby port into an Active port requires attaching it to the LAG, and updating its Actor

State to IN_SYNC. Once this occurs, both systems can move the newly active port into the distribution. Port 1 is relegated to Standby status, detached from the LAG, and its Actor and Partner States are updated accordingly.

Table 6-7
LAG state after distribution link failure, standard mode with actor priority

LAG port configuration			Port operational state			Actor state			Partner state		
Port	**Type**	**Priority**	**Oper**	**Select**	**Attach**	**Synch**	**Coll**	**Dist**	**Sync**	**Coll**	**Dist**
1	Dist	0x8000	Disabled	Standby							
2	Dist	0x8000	Enabled	Active	X	X	X	X	X	X	X
3	Prot	0x8000	Enabled	Standby					X		
4	Prot	0x4000	Enabled	Active	X	X	X	X	X	X	X

When port 1 returns to operationally enabled, Distribution port 1 must assume a Standby role since Distribution port 2 and Protection port 4 are currently the Active ports. As a result, port 1 is then not Attached, NOT_IN_SYNC and neither Collecting, nor Distributing.

If reversion is disabled on the LAG, then it stays in this state until either port 2 or port 4 becomes operationally disabled or deselected, allowing port 1, the Standby Distribution port, to resume its Active role. If reversion is enabled, then a reversion timer is started for each Distribution port when it returns to normal operation. When the reversion timer expires, the Distribution port is allowed to become Active, and forces the Protection port back into a Standby role.

The Example of initial LAG state, standard mode with partner priority table depicts the local LAG configuration and operational states of a LAG with two Distribution and two Protection ports operating in Standard protection mode when the partner system has a higher priority than the actor system. Both Distribution and Protection ports are considered to be Active, although the partner has chosen to use the Distribution ports. Since these ports are IN_SYNC on both systems, they are able to join the LAG distribution. The

Protection ports are OUT_OF_SYNC on the partner side, so they are unavailable for distributing traffic, although they are IN_SYNC on the actor side.

Table 6-8
Example of initial LAG state, standard mode with partner priority

LAG port configuration			Port operational state			Actor state			Partner state		
Port	**Type**	**Priority**	**Oper**	**Select**	**Attach**	**Synch**	**Coll**	**Dist**	**Sync**	**Coll**	**Dist**
1	Dist	0x8000	Enabled	Active	X	X	X	X	X	X	X
2	Dist	0x8000	Enabled	Active	X	X	X	X	X	X	X
3	Prot	0x8000	Enabled	Active	X	X					
4	Prot	0x4000	Enabled	Active	X	X					

The LAG state after distribution link failure, standard mode with actor priority table depicts the LAG after a link failure causes port 1 to become operationally disabled. In this example, the partner system has a higher priority and it chooses port 3 as the replacement link. Note that the partner system is able to apply its own selection logic and is not required to honor actor port priorities. In this case port 4 is actually higher priority from the actor's perspective.

Once its Partner State is updated to IN_SYNC, both systems can move the newly active port into the distribution. Port 1 continues to be an Active port, its Partner State is updated to OUT_OF_SYNC, and it is removed from the distribution.

Table 6-9
LAG state after distribution link failure, standard mode with actor priority

LAG port configuration			Port operational state			Actor state			Partner state		
Port	**Type**	**Priority**	**Oper**	**Select**	**Attach**	**Synch**	**Coll**	**Dist**	**Sync**	**Coll**	**Dist**
1	Dist	0x8000	Disabled	Active	X	X				X	X
2	Dist	0x8000	Enabled	Active	X	X	X	X	X	X	X
3	Prot	0x8000	Enabled	Active	X	X	X	X	X	X	X
4	Prot	0x4000	Enabled	Active	X	X					

After port 1 returns to operationally enabled, port 1, the Distribution port, becomes Active since the partner system has a higher priority. Since the Partner State of port 1 is NOT_IN_SYNC it does not join the distribution. Reversion settings on the actor system are also ignored and it is up to the

partner system to initiate replacing the Protection ports with the Distribution ports. When this occurs, the actor system sees a change in port Partner States and reacts accordingly.

Using the `aggregation show member` command, the state of each port in the aggregation group is displayed in the Type field. This field can help determine if a port is participating in the aggregation group or, in the case of a Protection port, whether a port is distributing traffic or just collecting traffic. Protection ports are indicated by a “P” next to the port number.

This table lists port states.

Table 6-10
Port states

Port state	Description
Added	The physical port has been administratively added to the logical aggregation port (for example, `aggregation add agg port` command).
Select	The physical port has been selected for this aggregation port by the protocol.
Agged	The physical port has been aggregated to this aggregation port by the protocol. This physical port now belongs the logical aggregation port.
Dist	The port is forwarding traffic over the logical aggregation port.

Minimum link aggregation

Minimum link aggregation provides a way to configure a minimum number of physical ports in an aggregation group that must be distributing traffic for the LAG to be considered operational. This feature can be used to determine whether a LAG is a desirable interface to use. For example, a LAG port that has four physical ports added to it, but only one of them is distributing traffic may not be a desirable interface.

The LAG port’s operational state is re-evaluated if one of these events occurs:

- The minimum link aggregation mode is enabled or disabled for a LAG
- A port is added to or removed from the LAG
- A member port of the LAG becomes operationally enabled or disabled
- The minimum link threshold changes
- The number of distributing ports in the LAG changes
- A peer node indicates an attribute change for any of its LAG member ports

When a minimum link aggregation is unsatisfied and its member ports come out of distribution, they are in the waiting/standby state. This means the port is not collecting or distributing. If the LAG is operating in LACP mode, the sync bit in the LACPDUs is turned off.

> ***Note:*** The operational state of the physical port does not change. If the LAG is operating in LACP mode, it continues to exchange LACPDUs.

Ports that are in waiting/standby state because the threshold was not met can be forced back into collecting or distributing either by disabling the minimum link aggregation mode or by modifying the minimum link threshold.

When the minimum link aggregation mode for a LAG is changed from enabled to disabled, the ports need to be re-selected to the link aggregation. If the partner port's state changed such that it is not in sync, the ports do not go into distribution.

Protection ports

If the LAG has protection ports added to it, the protection ports can take over if the primary ports go down. If the number of distributing protection ports combined with the distributing primary ports is equal to or exceeds the threshold, the LAG is still considered operationally enabled. This figure shows an example for LAG minimum link aggregation.

Figure 6-2
LAG minimum link aggregation example

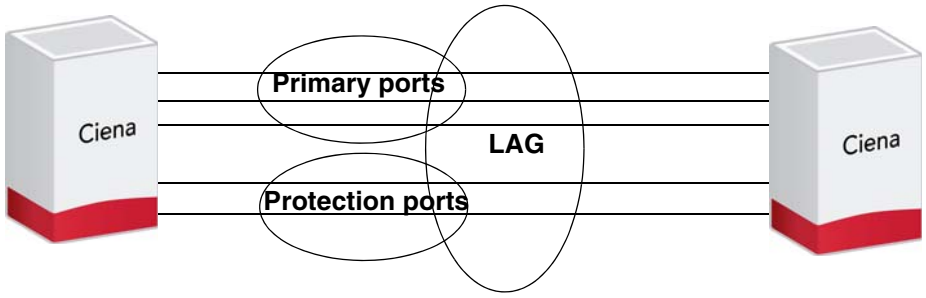

This example shows LAG minimum link aggregation behavior with protection ports. This LAG has three primary ports (5/1, 5/2 and 5/3) and two protection ports (5/4 and 5/5). The LAG has minimum link aggregation mode enabled and the link threshold is set to two. This means a minimum of two links must be distributing for the LAG to be operationally enabled. If 5/1 became disabled, protection port 5/4 starts distributing, resulting in three distributing links in the LAG. If both 5/2 and 5/3 become disabled, protection port 5/5 goes into distributing. The link threshold criterion is still met and the LAG is still operationally enabled. If one of the protection ports goes down, the LAG becomes operationally disabled since the link threshold is not being met.

It is recommended to configure minimum link aggregation on the LACP Master because minimum link aggregation overrides the LACP state machines by not allowing the operationally up ports to go to distribution if the minimum link threshold is not met. Minimum link aggregation forces those ports to be in waiting state and hence keeping the sync bit off. If minimum link aggregation is configured on the slave, a conflict between the master and slave results. This is because the master keeps the sync bit on and asks the slave to move the port to active/distribution while the slave tries to keep the port out of distribution. For minimum link aggregation to function properly, it must be configured on the master which has all the control.

Interoperability

If the remote device does not support minimum link aggregation, and the LAGs are operating in LACP mode when ports on the local device go out of distributing because the threshold is not being met, the corresponding LAG member ports on the remote device also go out of distributing and collecting and into the attached state. The local device that supports minimum link aggregation sends LACPDUs with the sync bit turned off so that the remote device knows how to keep the corresponding ports out of collecting and distributing.

If the LAGs are in manual mode with minimum link aggregation mode enabled, the threshold is not being met and the ports go out of distributing on the local end, the corresponding ports on the remote device continue collecting and distributing if the remote device does not support minimum link aggregation.

This chapter provides these procedures for link aggregation:

- "Configuring LACP between two devices" on page 6-26
- "Configuring LACP protection" on page 6-29
- "Configuring manual link aggregation with the IEEE 802.3ad MIB" on page 6-31
- "Enabling and disabling minimum link aggregation mode" on page 6-35
- "Setting the minimum link aggregation threshold" on page 6-36

Procedure 6-1
Configuring LACP between two devices

Configure LACP between two devices for load balancing and sharing, bandwidth aggregation, and link redundancy.

When configuring LACP, all links in the aggregation must either be up or down.

Note 1: For manual Link Aggregation, it is recommended that both devices be configured first, before the physical connections are made. Ciena does not recommend manual configuration of LACP parameters and keys unless the user has a firm understanding of the LACP standard.

Note 2: When naming Link Aggregation Groups, use a valid name string.

Note 3: When naming Link Aggregation Groups, do not use:

— a dash (-) symbol. The CLI interprets the dash as denoting a port range.

— a valid port number.

— an illegal character, such as slash (/).

Note 4: You cannot create an aggregation that uses one non-enhanced and one enhanced port on the 5150.

CAUTION

Possible Service Disruption

When changing the mode of a link aggregation from manual to LACP or from LACP to manual for a single link, a service disruption occurs on the LACP side until the modes match again. If connecting to the device using the management interface, the desired aggregation state must be modified on the “far end” device first or connectivity is lost unless a backup link exists.

For example, remote connections between Juniper devices are lost after disabling link aggregation. Juniper devices must be set to manual mode with the interfaces <agg group> aggregated-ether-options lacp command.

To minimize service disruption, Ciena recommends provisioning backup links before modifying the link aggregation mode.

Step	Action

1 Create a Link Aggregation Group (virtual port) on Device 1.

```
aggregation create agg <String[31]>
```

where

agg <String[31]>	is the aggregation port name.

2 Add ports to the Aggregation Group.

```
aggregation add agg <agg> port <Physical port list>
protection-port <Physical port list>
```

where

agg <agg>	is the aggregation port.
port <Physical port list>	is the list of ports to add to the aggregation.
protection-port <Physical port list>	is the list of protection ports to add to the aggregation.

3 Set the hashing method for the LAG.

```
aggregation set agg <agg> hash <enhanced | mac-addr-based
| ip-addr-based>
```

where

agg <agg>	is the aggregation port.
hash <enhanced | mac-addr-based | ip-addr-based>	is the hashing mode for known unicast frames.

4 Verify the configuration of the LAG.

```
aggregation show agg <agg> info mode
```

where

agg <agg>	is the aggregation priority.
info	displays aggregation information.
mode	displays aggregation mode.

5 Create a VLAN for the LAG

```
vlan create vlan <vlan> [name <String[31]>]
```

where

vlan <vlan>	is the VLAN ID.
name <String[31]>	is the name of the VLAN.

6 Add the aggregation to the VLAN

```
vlan add vlan <vlan> port <Port list>
```

where

vlan <vlan> is the VLAN ID.

port <Port list> is the list of ports to add to the VLAN.

7 Repeat the configuration for device 2.

—end—

Example

This example creates an aggregation group with ports 5,8,11, and 15 named group_1 in VLAN 10. Since the steps in configuration of both devices is the same for both link partners, only device 1 is shown.

```
> aggregation create agg group_1
> aggregation add agg group_1 port 5,8,11,15
> aggregation show agg group_1 info

+---------------------- Agg Port 2049 Info --------------------+
|          Parameter          |               Value              |
+-----------------------------+----------------------------------+
| LAG Port Name & ID          | group_1  0x0801(2049)            |
| Added  Total Ports          | 5 8 11 15                        |
|        Primary     Ports    |    5 8 11 15                     |
|        Protection  Ports    |    0  -                          |
| Selected Ports              | 5 8 11 15                        |
| Admin & Oper State          | Up   Up                          |
| Lacp Mode                   | LACP                             |
| Marker Delay                | 0                                |
| Marker Resp_All_Rcvd Count  | 0                                |
|        Time_Out Count       | 0                                |
| Ready Waiting               | None                             |
| Port Entry                  |                                  |
|    Aggregator Index         | 0x0801                           |
|    ACTOR Port MAC           | 00:02:A1:07:18:A0                |
|       Sys Prio & ID         | 0x8000 00:02:A1:07:18:A0         |
|       Admin & Oper Key      | 0x0801 0x2801                    |
|       Agg/Individual        | Aggregate                        |
|       Coll Max. Delay       | 0                                |
|    PARTNER Sys Prio & ID    | 0x8000 00:02:A1:15:57:40         |
|       Oper  Key             | 0x2801                           |
|       Coll Max Delay        | 0                                |
| Revert Time out             | 5000    (ms)                     |
| Revert Protection           | Off                              |
+-----------------------------+----------------------------------+
```

Note: The operational state value displays as Up as long as the aggregation contains at least one port in the distributing state. Otherwise, it is down.

```
> vlan create vlan 10
> vlan add vlan 10 port group_1
```

Procedure 6-2
Configuring LACP protection

Configure LACP protection for link redundancy.

Step	Action

1 Create a Link Aggregation Group (virtual port) on device 1.

```
aggregation create agg <String[31]>
```

where

agg <String[31]> is the aggregation port name.

2 Add ports to the Aggregation Group as regular Distribution ports.

```
aggregation add agg <agg> port <Physical port list>
protection-port <Physical port list>
```

where

agg <agg> is the aggregation port.

port <Physical port list> is the list of ports to add to the aggregation.

protection-port <Physical port list> is the list of protection ports to add to the aggregation.

3 Add additional ports to the Aggregation Group as Protection ports.

```
aggregation add agg <agg> port <Physical port list>
protection-port <Physical port list>
```

where

agg <agg> is the aggregation port.

port <Physical port list> is the list of ports to add to the aggregation.

protection-port <Physical port list> is the list of protection ports to add to the aggregation.

4 Enable or disable revert protection. The default value is off.

```
aggregation set agg <agg> revert-protection <on|off>
```

where

agg <String[31]> is the aggregation port name.

5 Set revert delay timer from.

```
aggregation set agg <agg> revert-delay <MILLISECONDS>
```

where

agg <agg>	is the aggregation port.
revert-delay <MILLISECONDS>	is the revert delay. Valid values are 0—60,000 ms. The default value is 5000 ms.

—end—

Example

Since the steps in configuration of both devices is the same for both link partners, only device 1 is shown in This example.

```
> aggregation create agg 111
> aggregation add agg 111 port 1,2,3
> aggregation add agg 111 protection-port 4,5,6
> aggregation set agg 111 revert-protection on
> aggregation set agg 111 revert-delay 10000
```

Procedure 6-3
Configuring manual link aggregation with the IEEE 802.3ad MIB

Configure manual link aggregation with the IEEE 802.3ad MIB.

You can set a physical port actor admin key without creating an aggregation port; Ciena recommends that you create an aggregation group and then set the physical port admin key for the aggregation group.

Note: Ciena does not recommend manual configuration of LACP parameters and keys unless the user has a firm understanding of the IEEE Standard for Local and metropolitan area networks Link Aggregation.

Step	Action

1 Create an aggregation group:

```
aggregation create {agg <String[31]>}
```

where

agg <String[31]> is the aggregation port name.

2 Set aggregation group attributes:

```
aggregation lacpmib set agg <agg> {admin-key <NUMBER:
0..4095>} {system-priority <NUMBER: 0-65535>}
```

where

agg <agg>	is the aggregation port name.
admin-key <NUMBER: 0..4095>	is the administrator key.
system-priority <NUMBER: 0..65535>	is the system priority.

3 Set physical port attributes:

```
aggregation lacpmib set port <port> {actor-admin-key
<NUMBER: 0..4095>} {actor-admin-state <NUMBER: 0..255>}
{actor-port-priority <NUMBER: 0..65535>} {actor-system-
priority <NUMBER: 0..65535>} {partner-admin-key <NUMBER:
0..4095>} {partner-admin-state <NUMBER: 0..255>}
```

```
{partner-port-priority <NUMBER: 0..65535>} {partner-
system-id <MAC address: XX:XX:XX:XX:XX:XX>} {partner-
system-priority <NUMBER: 0..65535>}
```

where

port <port>	is the physical port.
actor-admin-key <NUMBER: 0..4095>	is the actor administrator key.

where

actor-admin-state <NUMBER: 0..255>	is the actor administrator state: bit0_act bit1_timeout bit2_agg. This attribute is a string of 8 bits, corresponding to the administrative values of Actor_State as transmitted by the Actor in LACPDUs. • The first bit corresponds to bit 0 of Actor_State (LACP_Activity), which indicates the Activity control value. — Active LACP: 1 — Passive LACP: 0 • The second bit corresponds to bit 1 (LACP_Timeout), which indicates the Timeout control value. — Short Timeout:1 — Long Timeout: 0 • The third bit corresponds to bit 2 (Aggregation), which indicates whether the link is a candidate for aggregation or can only be operated as an individual link. — TRUE (the system considers this link able to be aggregated): 1 — FALSE (the system considers this link to be individual): 0 • The fourth bit corresponds to bit 3 (Synchronization), which indicates whether the system considers the link to be in the correct aggregation. — TRUE (IN_SYNC, that is, the link is allocated to the correct aggregation group, the group is associated with a compatible aggregator, and the identity of the LAG is consistent with the transmitted system ID and operational key information): 1 — FALSE (OUT_OF_SYNC, that is, the link is not in the correct aggregation): 0 • The fifth bit corresponds to bit 4 (Collecting). • The sixth bit corresponds to bit 5 (Distributing). • The seventh bit corresponds to bit 6 (Defaulted). • The eighth bit corresponds to bit 7 (Expired). These values allow administrative control over the values of LACP_Activity, LACP_Timeout and Aggregation.
actor-port-priority <NUMBER: 0..65535>	is the actor port priority.

where

actor-system-priority <NUMBER: 0..65535>	is the actor system priority.
partner-admin-key <NUMBER: 0..4095>	is the partner administrator key.
partner-admin-state <NUMBER: 0..255>	is the partner administrator state: bit0_act bit1_timeout bit2_agg.
partner-port-priority <NUMBER: 0..65535>	is the partner port priority.
partner-system-id <MAC address: XX:XX:XX:XX:XX:XX>	is the partner system ID (MAC).
partner-system-priority <NUMBER: 0..65535>	is the partner system priority.

4 Set the aggregation to manual mode:

```
aggregation set agg <agg-name> mode manual
```

—end—

Example

This example sets passive mode and long timeout (90 seconds):

```
aggregation lacpmib set port X actor-admin-state 4
```

This example sets active mode and long timeout (90 seconds):

```
aggregation lacpmib set port X actor-admin-state 5
```

This example sets passive mode and short timeout (3 seconds):

```
aggregation lacpmib set port X actor-admin-state 6
```

This example sets active mode and short timeout (3 seconds):

```
aggregation lacpmib set port X actor-admin-state 7
```

Procedure 6-4
Enabling and disabling minimum link aggregation mode

You can enable and disable the minimum link aggregation mode.

If the minimum link aggregation mode is enabled for an aggregation group, you can specify the minimum number of physical ports in the aggregation group that must be distributing traffic for the LAG to be considered operational.

Step	Action

To enable the minimum link aggregation mode

1 Enable the minimum link aggregation mode:

```
aggregation set agg <agg_name> min-link-aggregation on
```

where

agg <agg_name> is the aggregation port

To disable the minimum link aggregation mode

2 Disable the minimum link aggregation mode:

```
aggregation set agg <agg_name> min-link-aggregation off
```

where

agg <agg_name> is the aggregation port

—end—

Procedure 6-5
Setting the minimum link aggregation threshold

The threshold specifies the minimum number of physical ports in the aggregation group that must be distributing for the LAG to be considered operational. This threshold setting only takes effect if min-link-aggregation is enabled.

Step	Action
1	Set the minimum link aggregation threshold: `aggregation set agg <agg_name> min-link-aggregation-threshold {1-8}` where agg <agg_name> — is the aggregation port min-link-aggregation-threshold {1-8} — sets the minimum link threshold. The default is 1.

—end—

Link Layer Discovery Protocol (LLDP) configuration

The system software supports Link Layer Discovery Protocol (LLDP) as specified in the IEEE 802.1AB-2005 standard. Like Rapid Spanning Tree Protocol (RSTP), LLDP is a link-constrained, layer 2 protocol. This means that the exchange of LLDP messages only takes place between adjacent LLDP agents (devices) on the network, unless control frame tunneling is used to tunnel the Link Layer Discovery Protocol Data Units (LLDPDUs) through another device.

LLDP agents communicate basic management information with each other by exchanging LLDPDUs. The information about a device, and its immediate neighbors is then retrievable through the standard LLDP MIBs using SNMP. It is important to note that LLDP operates only on physical ports.

As illustrated in the Example of Architectural Relationship Between LLDP Agents figure, LLDP Agent A exchanges LLDPDUs with LLDP Agent B and Agent C, but not with Agent D since it is not directly connected to Agent A.

Figure 7-1
Example of Architectural Relationship Between LLDP Agents

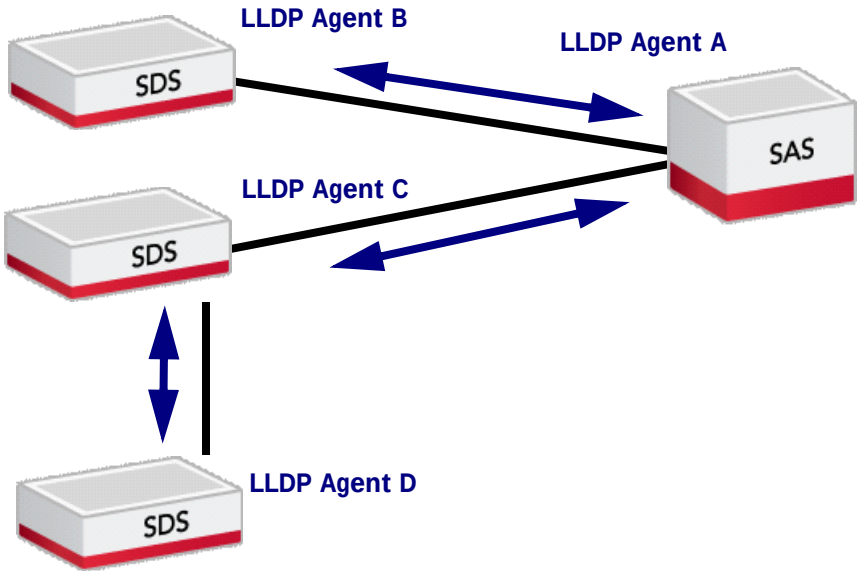

Note: Link Layer Discovery Protocol does not configure or control any traffic or devices on the network. Its primary role is to report discovered information to higher-layer management tools. It is not intended to act as a configuration protocol for remote systems nor as a mechanism to signal control information between ports.

LLDP TLVs

The LLDPDU is a Layer 2 packet that consists of an L2 source, destination, and an Ethertype field, and four or more Type Length and Value fields (TLVs) (see the LLDPDU Packet figure). The LLDP standard specifies nine common TLV fields, four of which are mandatory TLVs while the remaining five are optional TLVs to carry information for broadcasting sender information.

Figure 7-2
LLDPDU Packet

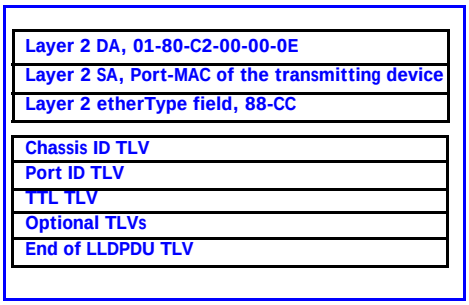

Periodic LLDPDUs are sent out at a user-defined interval; however, LLDPDUs are also sent whenever LLDP TLV data has changed. An SNMP notification is then generated after changed data is received from a neighboring device. Upon receipt of this notification the SNMP management application polls the LLDP MIB objects to determine what information has changed.

Common TLVs

The Common TLVs table lists LLDP TLVs and provides a description for each TLV. The system software supports the nine common TLVs specified by IEEE 802.1AB-2005.

Table 7-1
Common TLVs

TLV	Description
Common	
Chassis ID TLV	Identifies the chassis. This TLV contains the 802 LAN station associated with the transmitting LLDP agent and supports the “MAC Address” chassis ID subtype. This TLV is mandatory.
Port ID TLV	Identifies the port component that transmits the TLV. 39XX/51XX SDS devices support sending subtype “Interface Alias” and receiving subtypes “Interface Alias” and “MAC Address”. This TLV is mandatory.

Table 7-1
Common TLVs

TLV	Description
Time To Live (TTL) TLV	Indicates the number of seconds that the recipient LLDP agent is to regard the information associated with this LLDPDU to be valid. This TLV is mandatory.
Port Description TLV	Specifies the description for the port as an alphanumeric string. This TLV is optional.
System Name TLV	Specifies the assigned name as an alphanumeric string. This TLV is optional.
System Description TLV	Indicates the system description as an alphanumeric string. This TLV is optional.
System Capabilities TLV	Identifies the primary functions of the system and whether these primary functions are enabled. It supports the bridge capability. This TLV is optional.
Management Address TLV(s)	Identifies the IPv4 and IPv6 address associated with the local LLDP agent that can be used to reach higher layer entities to assist discovery by network management. This TLV is optional.
End of LLDPDU TLV	Defines the end of LLDPDU with all 0 values in two octets. This TLV is mandatory.
802.1 Organizationally Specific TLVs	
Port VLAN ID TLV	Allows a port to advertise its VLAN ID (PVID) that is associated with untagged or priority tagged frames.
Port and Protocol VLAN ID TLV	The system software stores values received from the remote partner for this TLV but does not transmit this TLV.
VLAN Name TLV	The system software stores values received from the remote partner for this TLV but does not transmit this TLV.

Table 7-1
Common TLVs

TLV	Description
Protocol ID TLV	Allows an 802 LAN station to advertise particular protocols that are accessible through the port. Currently, the system software advertises thesevprotocols: RSTP and 802.3ah OAM.
802.3 Organizationally Specific TLVs	
MAC/PHY Configuration/ Status TLV	Advertises auto negotiation support and status, PMD auto negotiation capability, and operational MAU type.
Power via MDI TLV	The system software stores values received from the remote partner for this TLV but does not transmit this TLV.
Link Aggregation TLV	Advertises whether the link is capable of being aggregated, whether the link is currently in an aggregation, and the port ID of the aggregation if it is in an aggregation. LLDP advertises Link Agg TLV in the PDU but Link Aggregation is not currently supported and is disabled.
Maximum Frame Size TLV	Advertises the maximum frame size capability.

Table 7-1
Common TLVs

TLV	Description
Ciena Organizationally Specific TLVs	
Unknown TLVs	The system software stores all unknown TLVs for all valid LLDP PDUs for retrieval.
LLDP Graceful Shut Down	When an LLDP agent is administratively disabled, it executes a graceful shutdown handshake. It sends out the last LLDP PDU that specifies the TTL as zero in TTL TLV. When its counterpart receives this LLDPDU, it cleans up all the LLDP information received.
Encoding of SyncE Info	This TLV conveys the following information: • SyncE capability of the port, which is one of — input only — output only — both — not supported • current SyncE configuration of the port, which is one of — input reference — output reference — both — not configured

Feature benefits

Using LLDP for network topology discovery offers these benefits to the network provider:

- Accurate network topology discovery and management
- Support of standard tools using SNMP
- Multi-vendor interoperability

Accurate network topology discovery and management

Many network management tools use layer 3 protocols to automate the discovery process and track topology configurations and changes. The use of layer 2 LLDP allows network management tools to quickly and accurately discover current network topologies and to track both intentional and unintentional topology changes. Network administrators and technicians can use the accurate and up-to-date topology information to more quickly diagnose network problems in the field.

Support of standard tools

Each LLDP agent maintains its own per-port table of information in a standard SNMP MIB. Updates are sent as needed to the closest connected device on the network. Users access the device and manage LLDP information through the command line interface (CLI). Users can display general information, for example, Chassis ID and port configuration. Users with superuser security privileges can configure LLDP options and manipulate TLVs. For example, a specific port can be configured to only transmit or receive LLDPDUs.

LLDPDUs are used to update standard LLDP Management information Databases (MIBs) allowing any standard SNMP application to monitor the information as it changes.

Multi-vendor interoperability

LLDP is a standards-based protocol. Using LLDP rather than a proprietary topology discovery protocol allows the network provider to interoperate with non-Ciena devices that also support LLDP.

Procedures for LLDP are:

- "Configuring LLDP" on page 7-8
- "Configuring TLV transmission" on page 7-10
- "Displaying LLDP neighbors" on page 7-12
- "Enabling and disabling SNMP notifications" on page 7-13

Procedure 7-1
Configuring LLDP

LLDP is enabled by default on all ports, but can be enabled or disabled for each port.

CAUTION
Performance during topology discovery
The default settings for LLDP must be sufficient to ensure proper topology discovery in most networks. Since Ethernet Services Manager (ESM) topology discovery relies on LLDP for higher performance topology discovery, care must be taken when modifying the LLDP configuration. In certain topologies, LLDP does not need to be forwarded, for example, when using non-LLDP devices such as hubs.

Step	Action

1 Display the current state of the port.

```
lldp show port <port>
```

where

<port> is the port to be disabled.

2 Disable the port.

```
lldp set port <port> disable
```

where

<port> is the port to be disabled.

3 To verify the configuration of the port, execute this command:

```
lldp show port <port> configuration
```

where

<port> is the port to be disabled.

—end—

Example

In the following configuration example, the user displays the state of port 10, then disables LLDP on this port. Note that by using the `lldp show port` command, information about the remote port (neighbor) can be seen in the LLDP Remote Port TLV section of the output. This displays the MAC address, port number, and other information of the connected remote port.

Display the current state of port 10.

```
lldp show port 10
+----------- LLDP Port Configuration -----------+
|       |        |  Basic   |  Org Specific  |      |
|       | Admin |TLV  Type| Dot1   |Dot3|    |      |
| Port | State |123456789|1234567|1234|    |Notif|
+------+-------+---------+-------+----+---+-----+
|10    | Tx-Rx |111111111|1  1111|1111|    | Off |
+------+-------+---------+-------+----+---+-----+
```

Set port 10 to disable.

```
lldp set port 10 disable
```

Verify the configuration of the port.

```
lldp show port 10 configuration
+----------- LLDP Port Configuration -----------+
|       |        |  Basic   |  Org Specific  |      |
|       | Admin |TLV  Type| Dot1   |Dot3|    |      |
| Port | State |123456789|1234567|1234|    |Notif|
+------+-------+---------+-------+----+---+-----+
|10    |Disable|111111111|1  1111|1111|    | Off |
+------+-------+---------+-------+----+---+-----+
```

Procedure 7-2
Configuring TLV transmission

Some optional TLVs can be excluded from transmission in the LLDP PDU by port.

Step	Action

1 Set TLV transmission parameters:

```
lldp tlvtx set port <port> {[mgmt-addr-local | mgmt-addr-local-ipv6] <on|off>} {[mgmt-addr-remote | mgmt-addr-remote-ipv6] <on|off>} port-descr <on|off>  system-cap <on|off>  system-descr <on|off>  system-name <on|off>
```

where

port <port>	is the port list.
mgmt-addr-local | mgmt-addr-local-ipv6 <on|off>	indicates whether to transmit the Local Management Address. The default value is on.
mgmt-addr-remote | mgmt-addr-remote ipv6 <on|off>	indicates whether to transmit the Remote Management Address. The default value is on.
port-descr <on|off>	indicates whether to transmit the port description. The default value is on.
system-cap <on|off>	indicates whether to transmit the system capabilities. The default value is on.
system-descr <on|off>	indicates whether to transmit the system description. The default value is on.
system-name <on|off>	indicates whether to transmit the system name. The default value is on.

2 Set TLV transmission dot1 parameters:

```
lldp tlvtx-dot1 set port <port> port-vlan-id <on|off> protocol-id-dot1x <on|off> protocol-id-lacp <on|off> protocol-id-oam <on|off>  protocol-id-stp <on|off>
```

where

port <port>	is the port list.
port-vlan-id <on|off>	indicates whether to transmit the port VLAN ID. The default value is on.
protocol-id-dot1x <on|off>	indicates whether to transmit the protocol ID xdot1x. The default value is on.

where

protocol-id-lacp <on|off>	indicates whether to transmit the protocol ID LACP. The default value is on.
protocol-id-oam <on|off>	indicates whether to transmit the protocol ID OAM. The default value is on.
protocol-id-stp <on|off>	indicates whether to transmit the protocol ID STP. The default value is on.

3 Set TLV transmission dot3 parameters:

```
lldp tlvtx-dot3 set port <port> link-agg <on|off> mac-phy-config <on|off> max-frame-size <on|off> power-via-mdi <on|off>
```

where

port <port>	is the port list.
link-agg <on|off>	indicates whether to transmit the link aggregation status. The default value is on.
mac-phy-config <on|off>	indicates whether to transmit the MAC Phy configuration. The default value is on.
max-frame-size <on|off>	indicates whether to transmit the maximum frame size. The default value is on.
power-via-mdi <on|off>	indicates whether to transmit the power via MDI status. The default value is on.

—end—

Procedure 7-3
Displaying LLDP neighbors

Display LLDP neighbors when you want to view this information about neighboring devices:

- local port
- remote port
- remote management address
- chassis identification
- system name

Use the information to troubleshoot connectivity issues to other network elements.

Step	Action
1	Display neighboring devices: `lldp show neighbors`

—end—

Example

This example shows sample output for the lldp show neighbors command.

```
lldp show neighbors
```

```
+-------------------------------------------------------------------------------------------------------------+
|                                                LLDP Neighbors                                               |
+-------------------------------------------------------------------------------------------------------------+
|  Local Addr: 10.5.5.16                                                                                      |
|  Local Addr: fde5:1d38:f018:1:202:5aff:fe01:ee1d                                                            |
| System Name: 5142-16                                                                                        |
| System Desc: Ciena 5142 Service Aggregation Switch                                                          |
+-------------+-------------------------------+-------------------------------------------------------------+
|Local        |                                             Remote                                            |
+-------------+-------------------------------+-------------------------------------------------------------+
|Port         |Port                           |Info                                                         |
+-------------+-------------------------------+-------------------------------------------------------------+
|1            |Ciena 3942 Service Delivel121212|  Chassis Id: 001094000001                                   |
|             |                               |   Mgmt Addr: 10.5.5.13                                      |
|             |                               |              fde5:1d38:f018:1:223:8aff:fef9:8ca0            |
|             |                               |              fe80::223:8aff:fef9:8cbf                       |
|             |                               | System Name: 3942-13                                        |
|             |                               | System Desc: Ciena 3942 Service Delivery Switch             |
|             |                               | SyncE Suppt: Not Supported                                  |
|             |                               | SyncE Confg: Not Configured                                 |
+-------------+-------------------------------+-------------------------------------------------------------+
|3            |1.3                            |  Chassis Id: 00238A624200                                   |
|             |                               |   Mgmt Addr: 10.5.5.18                                      |
|             |                               |              fde5:1d38:f018:1:223:8aff:fe62:4200            |
|             |                               |              fe80::223:8aff:fe62:423f                       |
|             |                               | System Name: 5150-18                                        |
|             |                               | System Desc: Ciena CN 5150 Service Aggregation Switch       |
|             |                               | SyncE Suppt: Input and Output                               |
|             |                               | SyncE Confg: Not Configured                                 |
+-------------+-------------------------------+-------------------------------------------------------------+
```

Procedure 7-4
Enabling and disabling SNMP notifications

LLDP supports the standard lldpRemTablesChange SNMP notification per port. This notification is disabled by default. When enabled, an SNMP notification is generated upon LLDP activity for the port when the value of the lldpStatsRemTableLastChangeTime changes.

You can

- enable SNMP notifications
- disable SNMP notifications

Step	Action

1 Enable notification of LLDP activity for a port:

```
lldp set port <port> notification on
```

where

port <port>	is the physical port to enable notification of LLDP activity for.
notification on	enables notification.

2 Disable notification of LLDP activity for a port:

```
lldp set port <port> notification off
```

where

port <port>	is the physical port to disable notification of LLDP activity for.
notification off	disables notification.

—end—

NETCONF/YANG configuration

This section describes Network Configuration Protocol (NETCONF)/YANG configuration. NETCONF provides a way to install, manipulate and delete the configuration of a 39XX/51XX switch. YANG is the data modeling language NETCONF uses. The YANG model schemas used by 39XX/51XX switches are described in *39XX/51XX Service Delivery, Aggregation and Virtualization Switches YANG Reference Model*.

NETCONF is a configuration protocol for network devices that also provides monitoring and notification as part of network management. The NETCONF server resides on the 39XX/51XX switch. The NETCONF client is a Network Management System (NMS) that sends configuration information to the server.

NETCONF runs alongside CLI and SNMP. It uses XML-based data encoding for the configuration data and for the protocol message. NETCONF can run 16 simultaneous sessions.

This figure shows the NETCONF conceptual layers.

Figure 8-1
NETCONF conceptual layers

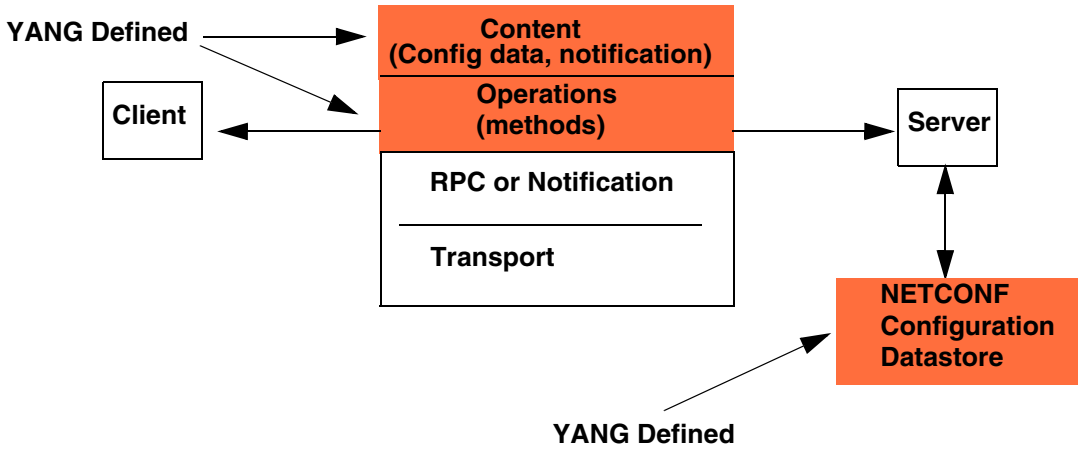

External NETCONF clients connect to a NETCONF server running on a 39XX/51XX switch to create, modify and delete configurations. NETCONF uses Remote Procedure Calls (RPCs) to communicate between the client and the server. The transport layer is responsible for authorization, encryption, and integrity of sessions between client and server.

NETCONF has four protocol commands as shown in this table.

Table 8-1
NETCONF protocol commands

Command name	Function
<get-config>	Returns configuration data only.
<get>	Returns configuration data and state data. Only configuration data is currently supported.
<edit-config>	Creates, modifies or deletes a configuration.
<copy-config>	Saves the complete SAOS configuration.

A single NETCONF <edit-config> request modifies one or more YANG objects, resulting in multiple SAOS calls. Multiple edits within a single <edit-config> request are serialized according to the NETCONF standard. Competing edits of the same object from multiple NETCONF clients can be executed in any order. The user is responsible for managing competing edits of the same YANG object. The <edit-config> is either applied or rejected.

Note: There is no way to way to roll back to an earlier <edit-config> request which has already succeeded in the previous operation.

A single NETCONF user account is required which has full access to all supported YANG models.

The NETCONF server process is disabled by default. Configuration synchronization during NETCONF startup or restart can take seconds to minutes, depending on the size of the configuration. During synchronization, the NETCONF login is blocked but the user can execute commands by means of the CLI.

Security

SSH provides authentication and session encryption for NETCONF. A NETCONF client can connect to SSH using the normal SSH listener port 22 or a special NETCONF-over-SSH listener port 830. Port 830 allows firewalls to easily identify and filter NETCONF-over-SSH traffic, but it can also let attackers identify NETCONF traffic. The NETCONF client requests the “netconf” subsystem after login to gain access to the NETCONF server.

Procedures

Procedures to support NETCONF/YANG configuration are:

- “Enabling and disabling NETCONF” on page 8-4
- “Configuring user access to NETCONF” on page 8-5
- “Configuring the SSH listener port for NETCONF” on page 8-6
- “Displaying NETCONF information” on page 8-7

Procedure 8-1
Enabling and disabling NETCONF

Enable the NETCONF server instance when the configuration of 39xx/51XX switches by means of NETCONF is required. Disable the NETCONF server instance when configuration by means of NETCONF is no longer required. By default, NETCONF is disabled.

If there is a system crash or failure, the NETCONF server process restarts automatically; however, if there are three failures within 15 minutes the NETCONF server process is disabled and must be restarted manually.

The NETCONF startup time depends on the configuration size.

Only one SAOS user name can have NETCONF access privileges.The NETCONF user has full read/write access to all YANG models, even if the user has read-only privileges in the CLI.

Step	Action

To enable the NETCONF server instance

1 Enable the SSH server:

```
ssh server enable
```

2 Add a user to NETCONF:

```
netconf add user <username>
```

where

user <username> is the specified SAOS user who has NETCONF privileges.

3 Enable the NETCONF server instance:

```
netconf enable
```

To disable the NETCONF server instance

4 Disable the NETCONF server instance:

```
netconf disable
```

—end—

Procedure 8-2
Configuring user access to NETCONF

Add NETCONF group access for a specified SAOS user to allow the SAOS user to perform NETCONF functions. Remove NETCONF group access for a specified SAOS user when the SAOS user is no longer required to perform NETCONF functions.

Only one SAOS user name can have NETCONF access privileges.

Note: You cannot change the NETCONF user while NETCONF is enabled.

Step	Action

To add NETCONF group access for a specified user

1 Add a specified SAOS user to NETCONF group:

```
netconf add user <username>
```

where

user <username> is the specified SAOS user who has NETCONF privileges.

To remove NETCONF group access for a specified user

2 Disable the NETCONF server instance:

```
netconf disable
```

3 Remove a specified SAOS user to NETCONF:

```
netconf remove user <username>
```

where

user <username> is the specified SAOS user who no longer has NETCONF privileges.

—end—

Procedure 8-3
Configuring the SSH listener port for NETCONF

Add the SSH listener port for NETCONF so that the NETCONF client can communicate with the NETCONF server. Remove the SSH listener port when the SSH listener port is no longer required.

Note: You cannot change the NETCONF listener port while NETCONF is enabled.

Step	Action

To add the SSH listener port for NETCONF

1 Add the SSH listener port for NETCONF:

```
netconf set listener-port <NUMBER: 22 to 65535>
```

where

listener-port <NUMBER: 22 to 65535>	sets the SSH listener port for NETCONF. Default port is 830.

To remove the SSH listener port for NETCONF

2 Disable the NETCONF server instance:

```
netconf disable
```

3 Remove the SSH listener port for NETCONF:

```
netconf unset listener-port
```

—end—

Procedure 8-4
Displaying NETCONF information

Display NETCONF operational status and configuration to verify operational status and configuration.

Step	Action
1	Display NETCONF operational status and configuration: `netconf show`

—end—

Example

This example shows output from the netconf show command when NETCONF is enabled.

```
> netconf show

+------------------------------ NETCONF ATTRIBUTES -----------------------
-----+
| Parameter               | Value                                          |
+-------------------------+------------------------------------------------
-----+
| Admin State             | Enabled                                        |
| Oper State              | Enabled                                        |
| Netconf Listener Port   | 830                                            |
| Netconf Users count     | 1                                              |
| Netconf User            | admin                                          |
+-------------------------+------------------------------------------------
-----+
```

This example shows output from the netconf show command when NETCONF is disabled.

```
> netconf show

+------------------------------ NETCONF ATTRIBUTES -----------------------
-----+
| Parameter               | Value                                          |
+-------------------------+------------------------------------------------
-----+
| Admin State             | Disabled                                       |
| Oper State              | Disabled                                       |
| Netconf Listener Port   | 830                                            |
| Netconf Users count     | 1                                              |
| Netconf User            | admin                                          |
+-------------------------+------------------------------------------------
-----+
```

39XX/51XX Service Delivery, Aggregation and Virtualization Switches

Base Configuration

SAOS 6.18
Publication: 009-3297-008
Document status: Standard
Revision A
Document release date: January 2019

CONTACT CIENA
For additional information, office locations, and phone numbers, please visit the Ciena web site at **www.ciena.com**